

Bayesian non-parametric inference for Λ -coalescents: posterior consistency and a parametric method

Jere Koskela
 j.j.koskela@warwick.ac.uk
 Mathematics Institute
 University of Warwick
 Coventry CV4 7AL
 UK

Paul A. Jenkins
 p.jenkins@warwick.ac.uk
 Department of Statistics
 University of Warwick
 Coventry CV4 7AL
 UK

Dario Spanò
 d.spano@warwick.ac.uk
 Department of Statistics
 University of Warwick
 Coventry CV4 7AL
 UK

August 25, 2016

Abstract

We investigate Bayesian non-parametric inference of the Λ -measure of Λ -coalescent processes with recurrent mutation, parametrised by probability measures on the unit interval. We give verifiable criteria on the prior for posterior consistency when observations form a time series, and prove that any non-trivial prior is inconsistent when all observations are contemporaneous. We then show that the likelihood given a data set of size $n \in \mathbb{N}$ is constant across Λ -measures whose leading $n-2$ moments agree, and focus on inferring truncated sequences of moments. We provide a large class of functionals which can be extremised using finite computation given a credible region of posterior truncated moment sequences, and a pseudo-marginal Metropolis-Hastings algorithm for sampling the posterior. Finally, we compare the efficiency of the exact and noisy pseudo-marginal algorithms with and without delayed acceptance acceleration using a simulation study.

1 Introduction

The Λ -coalescent family is a class of coalescent processes parametrised by probability measures on the unit interval, $\Lambda \in \mathcal{M}_1([0, 1])$, introduced by Donnelly and Kurtz [1999], Pitman [1999] and Sagitov [1999]. We focus in particular on the family of Λ -coalescents with recurrent, finite sites, finite alleles mutation, which we refer to simply as “ Λ -coalescents” throughout this paper. This class of processes will be introduced formally in Section 2.

In recent years, Λ -coalescents have gained prominence as population genetic models for species with a highly skewed family size distribution, particularly among marine species [Boom et al., 1994, Árnason, 2004, Eldon and Wakeley, 2006, Birkner and Blath, 2008], and have also been suggested as models of evolution under natural selection [Neher and Hallatschek, 2013]. The Λ -measure models skewness of the family size distribution. Thus, Λ represents an important confounding factor for inference from genetic data, as well as being a quantity of interest in its own right. Failure to properly account for uncertainty in Λ could lead to model misspecification

and incorrect inference. Consequently, likelihood-based inference for Λ -coalescents has also been an active area of research [Birkner and Blath, 2008, Birkner et al., 2011, Koskela et al., 2015a]. A review of Λ -coalescents and their use in population genetic inference can be found in [Birkner and Blath, 2009], and Steinrück et al. [2013] present a review of Beta-coalescent models for marine species. In this paper we make the following contributions:

1. We provide the first non-parametric analysis of inferring Λ -measures from genetic data. Our method is also the first instance of inferring Λ using the Bayesian paradigm.
2. We prove inconsistency of the posterior in full generality when data is contemporaneous, and consistency under conditions which can be verified in practice when the data set forms a time series.
3. We present an implementable parametrisation of the non-parametric inference problem by quotienting the infinite dimensional space $\mathcal{M}_1([0, 1])$ in a suitable, data-driven way. We believe this quotienting approach to have utility in infinite dimensional inference beyond the context of this work.
4. We implement a pseudo-marginal MCMC algorithm for sampling the posterior, and provide an illustrative simulation study which demonstrates the feasibility of the algorithm for inference.

The usual approach to inference is to focus on a parametric family of Λ -measures and infer the parameters from observations. Common choices of parametric family are $\Lambda(dr) = 2(2 + \psi^2)^{-1}\delta_0(dr) + \psi^2(2 + \psi^2)^{-1}\delta_\psi(dr)$ where $\psi \in (0, 1]$ [Eldon and Wakeley, 2006], $\Lambda = \text{Beta}(2 - \alpha, \alpha)$ where $\alpha \in (1, 2)$ [Birkner and Blath, 2008, Birkner et al., 2011] and $\Lambda(dr) = c\delta_0(dr) + 2^{-1}(1 - c)rdr$ where $c \in [0, 1]$ [Durrett and Schweinsberg, 2005]. We adopt the Bayesian non-parametric approach to circumvent restrictive, finite-dimensional parametrisations. We show in Theorem 1 that the posterior is inconsistent in the typical setting of sampling from a stationary population at a fixed time under any non-trivial prior (i.e. one which does not assign full mass to the truth), including parametric families. In Theorem 2 we adapt a result of [Koskela et al., 2015b] to provide verifiable criteria on the prior for posterior consistency when time series data is available. We also show in Section 5 that the popular Dirichlet process mixture model prior [Lo, 1984] satisfies these conditions. Recent advances in sequencing technology have made genetic time series data available [Drummond et al., 2002, Beaumont, 2003, Anderson, 2005, Drummond et al., 2005, Bollback et al., 2008, Minin et al., 2008, Malaspinas et al., 2012, Mathieson and McVean, 2013], and our results provide strong motivation for its continued use and development.

In Section 4 we make use of the fact that a sample of size $n \in \mathbb{N}$ carries information about $\Lambda \in \mathcal{M}_1([0, 1])$ only via its first $n - 2$ moments to make progress towards an implementable, non-parametric algorithm. This fact is implicit in well known sampling recursions for Λ -coalescent processes [Möhle, 2006, Birkner and Blath, 2008, 2009] and made explicit in Lemma 4. It has two important consequences. Firstly, it is natural to parametrise the inference problem with truncated moment sequences because any variation in the posterior between two Λ -measures whose first $n - 2$ moments coincide is due solely to the prior. Thus we obtain parametric inference algorithms requiring no discretisation or truncation, which are nevertheless as general as any non-parametric method. Secondly, while the conditions imposed on the prior by Theorem 2 are restrictive when viewed in $\mathcal{M}_1([0, 1])$ (e.g. they rule out all of the parametric families listed above), they are sufficiently mild that priors whose support contains an arbitrarily good approximation of the truncated moment sequence of any desired $\Lambda \in \mathcal{M}_1([0, 1])$ can be constructed. Posterior consistency of finite moment sequences is readily inherited from posterior consistency of the Λ -measure.

The truncated moment sequence approach can be thought of as regularising an underdetermined inference problem by identifying an appropriate, data-driven quotient space of parameters. This

is reminiscent of the method of *likelihood-informed subspaces* [Cui et al., 2014], which is a recently developed tool for efficient MCMC for inverse problems. The difference between existing subspace approaches and our quotient space is that existing work has focused on projections onto finite dimensional subspaces which preserve “most” of the likelihood in some sense, while our method captures the likelihood function exactly. Hence we believe our work will have wider utility beyond the Λ -coalescent setting.

We show in Theorem 3 that identifying $n - 2$ moments of Λ does not correspond to identifying smaller regions in $\mathcal{M}_1([0, 1])$ with increasing n , when measured in either total variation or Kullback-Leibler divergence. Hence the straightforward approach of computing a maximum a posteriori moment sequence, identifying a candidate $\Lambda^* \in \mathcal{M}_1([0, 1])$ with that moment sequence and treating Λ^* as a point estimator is inappropriate. Instead we use results by Winkler [1988] to provide a broad class of functionals of Λ which can be maximised or minimised over credible regions of the posterior. As a simple example, it is often of great interest whether the Kingman coalescent [Kingman, 1982], $\Lambda = \delta_0$, provides an adequate model for observed genetic data. Any parametric family containing δ_0 as a special case could be used to assess the Kingman assumption, but such an approach requires justification of the parametric family. It is also possible that δ_0 is a good model within the family, but a poor fit in some broader one. Our method avoids both of these problems.

Finally, we provide a pseudo-marginal Metropolis-Hastings algorithm [Beaumont, 2003, Andrieu and Roberts, 2009] for sampling truncated moment sequences from the posterior distribution. We compare the performance of the standard pseudo-marginal algorithm, the noisy pseudo-marginal algorithm, as well as delayed acceptance accelerated versions of both algorithms [Christen and Fox, 2005] in order to reduce the number of expensive likelihood estimations and improve the computational feasibility of our inference. We find that the noisy algorithm does not improve upon standard pseudo-marginal inference in this setting, especially given that twice as many likelihood evaluations are required. In contrast, an off-the-shelf delayed acceptance step dramatically speeds up computations. These results are illustrated with a simulation study comparing the Kingman hypothesis, $\Lambda = \delta_0$, based on simulated data from both the Kingman and Bolthausen-Sznitman ($\Lambda = U(0, 1)$) coalescents. The Kingman coalescent is the classical null model of neutral evolution, while the Bolthausen-Sznitman coalescent has recently emerged as an alternative model in the presence of selection [Schweinsberg, 2015] or population expansion into uninhabited territory [Berestycki et al., 2013]. In particular, the Bolthausen-Sznitman coalescent has been suggested as a model for the genetic ancestry of microbial populations such as influenza or HIV [Neher and Hallatschek, 2013].

The rest of the paper is laid out as follows. Section 2 provides an introduction to Λ -coalescents, Λ -Fleming-Viot jump-diffusions and a duality relation connecting the two. In Section 3 we state and prove our consistency results. Section 4 presents our parametrisation via moments and shows that it preserves all information in the data. Section 5 contains example families of priors which satisfy both consistency and tractable push-forward priors on moment sequences. Section 6 contains our results on inferring functionals of Λ based on finitely many moments. In Section 7 we present the pseudo-marginal Metropolis-Hastings algorithm for sampling posterior distributions of moment sequences, and an accompanying simulation study which demonstrates the practicality of the method. Section 8 concludes with a discussion.

2 Preliminaries

Let $[d] := \{1, \dots, d\}$ represent d genetic types (e.g. $d = 4^l$ for l loci of DNA), $M = (M_{ij})_{i,j=1}^d$ be a stochastic matrix specifying mutation probabilities between types, $\theta > 0$ be the mutation rate and $\Lambda \in \mathcal{M}_1([0, 1])$ denote the Λ -measure. Here and throughout we assume M has a unique

stationary distribution $m \in \mathcal{M}_1([d])$.

Let $\mathbf{X} := \{\mathbf{X}_t\}_{t \geq 0}$ denote the Λ -Fleming-Viot process with recurrent, finite sites, finite alleles mutation: a jump-diffusion on the d -dimensional probability simplex $\mathcal{S}_d := \{\mathbf{x} \in [0, 1]^d : \sum_{i=1}^d x_i = 1\}$ with generator

$$\begin{aligned} \mathcal{G}^\Lambda f(\mathbf{x}) = & \frac{\Lambda(\{0\})}{2} \sum_{i,j=1}^d x_i(\delta_{ij} - x_j) \frac{\partial^2}{\partial x_i \partial x_j} f(\mathbf{x}) + \theta \sum_{i,j=1}^d x_j(M_{ji} - \delta_{ij}) \frac{\partial}{\partial x_i} f(\mathbf{x}) \\ & + \sum_{i=1}^d x_i \int_{(0,1]} [f((1-r)\mathbf{x} + r\mathbf{e}_i) - f(\mathbf{x})] r^{-2} \Lambda(dr) \end{aligned} \quad (1)$$

acting on functions $f \in C^2(\mathcal{S}_d)$. This process is a model for the distribution of genetic types $[d]$ in a large population undergoing recurrent mutation and random mating with high-fecundity reproduction events, in which a single individual becomes ancestral to a non-trivial fraction $r \in (0, 1]$ of the whole population. As with Λ -coalescents, we will abbreviate “ Λ -Fleming-Viot process with recurrent, finite sites, finite alleles mutation” to just “ Λ -Fleming-Viot process” throughout this paper.

The first term on the right hand side (henceforth R.H.S.) of (1) models diffusion of the allele frequencies due to random mating, or genetic drift in the terminology of population genetics. The second term models mutation, and the third (jump) term models high-fecundity reproductive events. Without the jump term, i.e. when $\Lambda = \delta_0$, $\{\mathbf{X}_t\}_{t \geq 0}$ reduces to the classical d -dimensional Wright-Fisher diffusion with recurrent mutation. We denote the law of a Λ -Fleming-Viot process with initial condition $\mathbf{x} \in \mathcal{S}_d$ by $\mathbb{P}_{\mathbf{x}}^\Lambda(\cdot)$, and expectation with respect to this law by $\mathbb{E}_{\mathbf{x}}^\Lambda[\cdot]$. We suppress dependence on initial conditions whenever the stationary process is meant. For bounded $f : \mathcal{S}_d \mapsto \mathbb{R}$, let $P_t^\Lambda f(\mathbf{x}) := \mathbb{E}_{\mathbf{x}}^\Lambda[f(\mathbf{X}_t)]$ be the associated transition semigroup, $p_t^\Lambda(\mathbf{x}, \mathbf{y})$ be the transition density and $\pi^\Lambda(\mathbf{x})$ be the corresponding stationary density on \mathcal{S}_d , assumed unique. The transition semigroup is Feller for any $\Lambda \in \mathcal{M}_1([0, 1])$ [Bertoin and Le Gall, 2003], and all densities are assumed to exist with respect to the d -dimensional Lebesgue measure $d\mathbf{x}$.

A realisation of a Λ -Fleming-Viot process specifies the relative frequencies of genetic types in an infinite population across time. It is natural to imagine sampling $n \in \mathbb{N}$ individuals from the population at a fixed time, and ask about the ancestral tree connecting the sampled individuals. This ancestral tree is a random object, and is described by the Λ -coalescent, denoted by $\Pi := \{\Pi_t\}_{t \geq 0}$, taking values in \mathcal{P}_n^d , the set of $[d]$ -labelled partitions of $[n]$ [Donnelly and Kurtz, 1999, Section 5]. The process is started from the unlabelled partition $\psi_n := \{\{1\}, \{2\}, \dots, \{n\}\}$, and propagates backwards in time from the point of view of the reproductive evolution of the population. When the process has $p \in \mathbb{N}$ blocks, any $2 \leq k \leq p$ of them merge at rate

$$\lambda_{p,k} := \Lambda(\{0\}) \mathbb{1}_{\{2\}}(k) + \int_{(0,1]} (1-r)^{p-k} r^{k-2} \Lambda(dr),$$

where $\mathbb{1}_A(k)$ is the indicator function of the set A . We will abbreviate $\lambda_{k,k} \equiv \lambda_k$ throughout the paper. Once the process has merged into a single block, known as the most recent common ancestor (MRCA), an ancestral type is sampled from an initial law $\nu \in \mathcal{M}_1([d])$. We focus on the case of a stationary population, in which case $\nu = m$. This type is inherited along the branches of the Λ -coalescent tree, with mutations occurring at rate θ and mutant types sampled from M . The result is a random labelling of partitions for all times $0 \leq t \leq T$, where T denotes the hitting time of the MRCA. Let \mathbf{P}_n^Λ denote the law of Π started from ψ_n , and E_n^Λ denote the corresponding expectation. For the entirety of the paper we assume θ and M are known, and focus on inferring Λ . This assumption is crucial for the proof of consistency of Bayesian inference from time series data (Theorem 2 in Section 3), but not needed for correctness of the finitely many moments parametrisation in Section 4, or of the algorithms in Section 7.

The following relationship between Λ -Fleming-Viot processes and corresponding Λ -coalescents is classical [Bertoin and Le Gall, 2003], and will be useful in the (in)consistency proofs in the next section:

$$\mathbb{E}^\Lambda \left[\prod_{i=1}^d X_t(i)^{n_i} \right] = E_n^\Lambda \left[\prod_{i=1}^d m(i)^{|\Pi_t(i)|} \right], \quad (2)$$

where n_i denotes the number of observed individuals with label $i \in [d]$ sampled i.i.d. from the random measure \mathbf{X}_t , and $|\Pi_t(i)|$ denotes the number of blocks in partition Π_t with label $i \in [d]$. Formula (2) is an example of so called *moment duality* between stochastic processes (see e.g. [Möhle, 1999], and references therein for details): Intuitively, (2) states that the law of the type frequencies of Λ -coalescent leaves coincides with a multinomial sample from a random measure drawn from the corresponding Λ -Fleming-Viot process.

3 Posterior consistency

Let $Q \in \mathcal{M}_1(\mathcal{M}_1([0, 1]))$ be a prior distribution for Λ , and $\mathbf{n} = (n_1, \dots, n_d) \in \mathbb{N}^d$ denote observed type frequencies of $n := \sum_{i=1}^d n_i$ $[d]$ -labelled lineages generated by the Λ -coalescent. For Borel sets $A \in \mathcal{B}(\mathcal{M}_1([0, 1]))$, define the posterior as

$$Q(A|\mathbf{n}) = \frac{\int_A \mathbf{P}_n^\Lambda(\mathbf{n}) Q(d\Lambda)}{\int_{\mathcal{M}_1([0, 1])} \mathbf{P}_n^\Lambda(\mathbf{n}) Q(d\Lambda)}.$$

Informally, posterior consistency holds when $Q(\cdot|\mathbf{n})$ concentrates on a neighbourhood of the $\Lambda_0 \in \mathcal{M}_1([0, 1])$ which generated \mathbf{n} as $n \rightarrow \infty$. This is a natural requirement for statistical inference as it ensures the truth can be learned from a sufficient amount of data. For an overview of Bayesian non-parametric statistics, the reader is directed at [Hjort et al., 2010] and references therein.

It is well known that non-parametric posterior consistency is highly sensitive to the details of the topology defining the neighbourhood system as well as the mode of convergence [Diaconis and Freedman, 1986]. This will also be the case for our consistency result for time series data, Theorem 2. In contrast, the inconsistency result for contemporaneous observations, Theorem 1, is very universal. Hence we postpone specification of these details until after Theorem 1.

Remark 1. In Theorem 1 below, we give a formal statement of inconsistency from the point of view of Bayesian estimators, which are our interest in this paper. Elements of the proof of Theorem 1 will also be useful in the proving Theorem 2 later in the paper. However, we emphasize that the negative result also holds for frequentist estimators based on contemporaneous data. Essentially the same argument used to prove Theorem 1 shows that the limiting likelihood

$$\lim_{n \rightarrow \infty} \mathbb{E}^\Lambda [q(\mathbf{n}|\mathbf{X})] = \pi^\Lambda(\mathbf{x})$$

is positive for any observation \mathbf{x} and any Λ , at least provided $\{\pi^\Lambda\}_\Lambda$ is a bounded family of stationary densities so that the Dominated Convergence Theorem holds without the regularising effect of the prior. Hence, estimators cannot converge to the true Λ -measure generating the data. Ultimately, the problem is that the d -dimensional vector corresponding to one exact draw from a stationary Λ -Fleming-Viot process cannot be expected to uniquely identify an infinite dimensional object. Similar problems of identifiability also occur in estimation of lifetime distributions in branching processes [Höpfner et al., 2002, Hoffmann and Olivier, 2016], for which identifiability issues can be overcome by letting the length of the observation window grow to infinity, analogously to Theorem 2 below.

Theorem 1. Let $\mathbf{n} \in \mathbb{N}^d$ denote the observed type frequencies in a sample of size $n \in \mathbb{N}$ generated by a Λ -coalescent started from ψ_n at a fixed time, and let $\mathbf{x} := \lim_{n \rightarrow \infty} \frac{\mathbf{n}}{n} \in \mathcal{M}_1([d])$ denote the limiting observed relative type frequencies. Then the limiting posterior is given by

$$\lim_{n \rightarrow \infty} Q(A|\mathbf{n}) = \frac{\int_A \pi^\Lambda(\mathbf{x}) Q(d\Lambda)}{\int_{\mathcal{M}_1([0,1])} \pi^\Lambda(\mathbf{x}) Q(d\Lambda)}.$$

In particular, the R.H.S. is positive for any $A \in \mathcal{B}(\mathcal{M}_1([0,1]))$ which intersects the support of Q , regardless of the $\Lambda \in \mathcal{M}_1([0,1])$ generating the data.

Proof. Conditioning on the ancestral tree of the observed sample give the following representation for the posterior:

$$Q(A|\mathbf{n}) = \frac{\int_A \mathbf{P}_n^\Lambda(\mathbf{n}) Q(d\Lambda)}{\int_{\mathcal{M}_1([0,1])} \mathbf{P}_n^\Lambda(\mathbf{n}) Q(d\Lambda)} = \frac{\int_A E_n^\Lambda [\mathbf{1}_{\{\mathbf{n}\}}(\Pi_0)] Q(d\Lambda)}{\int_{\mathcal{M}_1([0,1])} E_n^\Lambda [\mathbf{1}_{\{\mathbf{n}\}}(\Pi_0)] Q(d\Lambda)}.$$

Using (2) we can write

$$Q(A|\mathbf{n}) = \frac{\int_A \mathbb{E}^\Lambda [q(\mathbf{n}|\mathbf{X})] Q(d\Lambda)}{\int_{\mathcal{M}_1([0,1])} \mathbb{E}^\Lambda [q(\mathbf{n}|\mathbf{X})] Q(d\Lambda)},$$

where $q(\mathbf{n}|\mathbf{X}) := \binom{n}{n_1, \dots, n_d} \prod_{i=1}^d X_i^{n_i}$ is the multinomial sampling probability. We will show the requisite convergence of the numerator and denominator separately, and the result will follow by the algebra of limits. We begin by considering the numerator.

By Fubini's theorem

$$\int_A \mathbb{E}^\Lambda [q(\mathbf{n}|\mathbf{X})] Q(d\Lambda) = \int_{\mathcal{S}_d} q(\mathbf{n}|\mathbf{y}) \int_A \pi^\Lambda(\mathbf{y}) Q(d\Lambda) d\mathbf{y} =: \int_{\mathcal{S}_d} q(\mathbf{n}|\mathbf{y}) F_{Q;A}(\mathbf{y}) d\mathbf{y},$$

where $F_{Q;A}(\mathbf{y})$ is a sub-probability density on \mathcal{S}_d since it is a mixture of probability densities. Hence $F_{Q;A}(\mathbf{y}) d\mathbf{y}$ defines a finite measure on \mathcal{S}_d , and $q(\mathbf{n}|\mathbf{y}) \leq 1$ so that by the Dominated Convergence theorem

$$\lim_{n \rightarrow \infty} \int_{\mathcal{S}_d} q(\mathbf{n}|\mathbf{y}) F_{Q;A}(\mathbf{y}) d\mathbf{y} = \int_{\mathcal{S}_d} \lim_{n \rightarrow \infty} q(\mathbf{n}|\mathbf{y}) F_{Q;A}(\mathbf{y}) d\mathbf{y}.$$

By the Law of Large Numbers $\mathbf{n} \sim \lfloor n\mathbf{x} \rfloor$ so that $q(\mathbf{n}|\mathbf{y}) \rightarrow q(\lfloor n\mathbf{x} \rfloor|\mathbf{y})$, and by Stirling's formula

$$q(\lfloor n\mathbf{x} \rfloor|\mathbf{y}) \sim \prod_{i=1}^d \left(\frac{y_i}{x_i} \right)^{nx_i},$$

or

$$\log(q(\lfloor n\mathbf{x} \rfloor|\mathbf{y})) \sim n \sum_{i=1}^d x_i \log \left(\frac{y_i}{x_i} \right) = -n \sum_{i=1}^d x_i \log \left(\frac{x_i}{y_i} \right) = -n \text{KL}(\mathbf{x}, \mathbf{y}),$$

where $\text{KL}(\mathbf{x}, \mathbf{y})$ denotes the Kullback-Leibler divergence between the probability mass functions \mathbf{x} and \mathbf{y} . By Gibbs' inequality $\text{KL}(\mathbf{x}, \mathbf{y}) \geq 0$ and $\text{KL}(\mathbf{x}, \mathbf{y}) = 0$ if and only if $\mathbf{x} = \mathbf{y}$, so that $q(\mathbf{n}|\cdot) \rightarrow \delta_{\mathbf{x}}(\cdot)$ almost surely. Hence

$$\int_{\mathcal{S}_d} \lim_{n \rightarrow \infty} q(\mathbf{n}|\mathbf{y}) F_{Q;A}(\mathbf{y}) d\mathbf{y} = F_{Q;A}(\mathbf{x}) = \int_A \pi^\Lambda(\mathbf{x}) Q(d\Lambda),$$

as required. The argument for the denominator is identical after substituting $\mathcal{M}_1([0,1])$ for the domain of integration A . \square

Remark 2. There is an apparent contradiction between the negative conclusion of Theorem 1 and recent positive results [Spence et al., 2016, Theorems 2, 3, 4 and 5] showing that Λ -measures can often be identified from their site frequency spectra. The contradiction is resolved by noting that Spence et al. [2016] work directly with the expected site frequency spectrum, thereby sidestepping both the randomness of the ancestral tree and the randomness of the mutation process given the tree. Numerical investigations by Spence et al. [2016] show that their method is unreliable unless a modest number (10-100) of independent realisations of ancestral trees is available. Independent trees cannot be sampled from populations whose ancestry is described by any non-Kingman Λ -coalescent, even in the idealised scenario of an infinitely long genome in the presence of recombination. However, as noted by [Spence et al., 2016], in practice the decay of correlations with increasing genome length is determined by the prelimiting model of evolution, and not necessarily the limiting Λ -coalescent. For example, the selective sweep model of Durrett and Schweinsberg [2005] can allow for asymptotically independent trees across a genome in the presence of multiple mergers for some combinations of parameters, in which case the identifiability results of Spence et al. [2016] hold.

The following example is an extension of a result by Der and Plotkin [2014], and demonstrates that the lack of consistency can have dramatic consequences for statistical identifiability even in the case of very simple priors.

Example 1. Consider $d = 2$, $M_{ij} = 1/2$ for $i, j \in \{1, 2\}$, and set $Q(d\Lambda) = \frac{1}{2}\delta_{\delta_0}(d\Lambda) + \frac{1}{2}\delta_{\delta_1}(d\Lambda)$. The stationary law $\pi^\Lambda(x)$ is known in the parent-independent, two-allele case for both of these atoms [Der and Plotkin, 2014]:

$$\begin{aligned}\pi^{\delta_0}(x) &= \frac{\Gamma(2\theta)}{[\Gamma(\theta)]^2} x^{\theta-1} (1-x)^{\theta-1} \\ \pi^{\delta_1}(x) &= \frac{1}{\theta} |1-2x|^{\frac{1-\theta}{\theta}},\end{aligned}$$

so the expected limiting posterior probabilities can be computed assuming either data-generating measure. These are listed in Table 1 for some candidate values of θ , while Figure 1 depicts limiting posterior probabilities as functions of the observed allele frequencies. The extreme sensitivity of the posterior probabilities in Figure 1 is akin to the “Bayesian brittleness” investigated by Owhadi et al. [2015], resulting in inferences which are not robust to small changes in the observed allele frequencies, prior probabilities or latent parameters.

θ	$\mathbb{E}^{\delta_0}[Q(\delta_0 X)]$	$\mathbb{E}^{\delta_1}[Q(\delta_0 X)]$
0.04	0.84	0.16
0.1	0.73	0.27
0.5	0.54	0.46
1	0.50	0.50
5	0.65	0.35
10	0.75	0.25
17	0.82	0.18

Table 1: Expected posterior probabilities given an infinite number of contemporaneous observations in the parent-independent, two-allele model.

The fact that $\pi^{\delta_0}(x) = \pi^{\delta_1}(x)$ when $\theta = 1$ in Example 1 was pointed out by Der and Plotkin [2014] as proof of the fact that Λ -measures cannot in general be uniquely identified from independent draws from $\pi^\Lambda(x)$. Our calculations illustrate that inference suffers from low power even when $\theta \neq 1$ if all observations are contemporaneous.

The inconsistency result of Theorem 1 holds for essentially arbitrary priors. Our next aim is to show that the posterior can be consistent when the data set is a time series of increasing




Figure 1: Limiting posterior probabilities as functions of the observed allele frequency x for $\theta = 0.1$ and 10 in the parent-independent, two-allele model. $Q(\delta_0|x)$ is dotted, $Q(\delta_1|x)$ is dash-dotted, $\pi^{\delta_0}(x)$ is solid and $\pi^{\delta_1}(x)$ is dashed. Note the extreme sensitivity of the posterior to the observed allele frequencies near $x = 0.5$ when $\theta = 10$ and near $x = 0$ or 1 when $\theta = 0.1$.

length. This does not contradict the unidentifiability claim of Der and Plotkin [2014], because the authors only consider independent draws from π^Λ . In contrast, in our setting it is information about transition densities $p_\Delta^\Lambda(\mathbf{x}, \mathbf{y})$ which facilitates posterior consistency.

We begin by defining the topology and weak posterior consistency following the set up of van der Meulen and van Zanten [2013], who considered similar time series data for one dimensional diffusions. In addition to topological details, posterior consistency is highly sensitive to the support of the prior, which should not exclude the truth. Often this is guaranteed by insisting that the prior places positive mass on all neighbourhoods of the truth, typically measured in terms of Kullback-Leibler divergence. In our setting such a support condition is provided by (3) below.

Definition 1. Fix $\eta > 0$ and let \mathcal{D}_η be a collection of Lebesgue probability densities on $[\eta, 1]$ satisfying $\inf_{r \in [\eta, 1]} \phi(r) > 0$ and $\sup_{r \in [\eta, 1]} \phi(r) < \infty$ for each $\phi \in \mathcal{D}_\eta$. We assume that $\Lambda(dr) = \phi(r)dr$ for $\phi \in \mathcal{D}_\eta$, and denote the data generating density by ϕ_0 .

Restricting the support of ϕ to $[\eta, 1]$ ensures that the Λ -coalescent can have no Kingman component, and that the Λ -Fleming-Viot process is a compound Poisson process with drift. Furthermore, most previously studied parametric families of Λ -measures are ruled out, including all those mentioned in Section 1. However, we will see in Section 4 that the prior can be chosen to satisfy the conditions of Definition 1 and place mass arbitrarily close to any desired Λ -measure, or family of Λ -measures, in a way we will make precise in Example 2.

Lemma 1. For any pair $\phi \neq \phi' \in \mathcal{D}_\eta$ and any $\Delta > 0$ there exists $\mathbf{x} \in \mathcal{S}_d$ and $f \in D(\mathcal{G}^\phi) \subset C(\mathcal{S}_d)$, the domain of the generator, such that $P_\Delta^{\phi'} f(\mathbf{x}) \neq P_\Delta^\phi f(\mathbf{x})$, where $C(\mathcal{S}_d)$ denotes the set of continuous functions on \mathcal{S}_d . In particular, identifying P_Δ^ϕ is equivalent to identifying ϕ .

Proof. This follows immediately from Lemma 1 of [Koskela et al., 2015b]. \square

Definition 2. Fix a sampling interval $\Delta > 0$ and a finite Borel measure ν on \mathcal{S}_d placing positive mass in all non-empty open sets. A weak topology on \mathcal{D}_η is generated by open sets of the form

$$U_{f, \varepsilon}^\phi := \{\phi' : \|P_\Delta^{\phi'} f(\mathbf{x}) - P_\Delta^\phi f(\mathbf{x})\|_{1, \nu} < \varepsilon\},$$

for any $\phi \in \mathcal{D}_\eta$, $\varepsilon > 0$ and $f \in C(\mathcal{S}_d)$, where $\|\cdot\|_{p, \nu}$ is the $L^p(\mathcal{S}_d, \nu)$ -norm. The Lebesgue measure is meant whenever no measure is specified.

Lemma 2 of [Koskela et al., 2015b] yields that the topology generated by $U_{f,\varepsilon}^\phi$ is Hausdorff, and hence separates points.

Definition 3. Let $\mathbf{n}_0, \dots, \mathbf{n}_m$ denote $m + 1$ samples observed at times $0, \Delta, \dots, \Delta m$ from a stationary Λ -coalescent, with each sample being of size $n \in \mathbb{N}$. See e.g. [Beaumont, 2003] for details of how temporally structured samples can be generated. *Weak posterior consistency* holds if $Q(U_{\phi_0}^c | \mathbf{n}_0, \dots, \mathbf{n}_m) \rightarrow 0$ \mathbb{P}^{ϕ_0} -a.s. as $n, m \rightarrow \infty$, where U_{ϕ_0} is any open neighbourhood of $\phi_0 \in \mathcal{D}_\eta$.

Theorem 2. Let $\mathbf{n}_0, \dots, \mathbf{n}_m$ be as in Definition 3 and $\mathbf{x}_0, \dots, \mathbf{x}_m$ denote the observed limiting type frequencies as $n \rightarrow \infty$, i.e. $\mathbf{x}_i = \lim_{n \rightarrow \infty} \mathbf{n}_i/n$. Suppose that the prior Q is supported on a set \mathcal{D}_η which satisfies the conditions in Definition 1, and for any $\varepsilon > 0$ and any $\phi_0 \in \mathcal{D}_\eta$ suppose that

$$Q\left(\phi \in \mathcal{D}_\eta : \int_\eta^1 \left\{ \left| \log\left(\frac{\phi_0(r)}{\phi(r)}\right) \right| + \left| \frac{\phi_0(r)}{\phi(r)} - 1 \right| \right\} r^{-2} \phi_0(r) dr < \varepsilon\right) > 0. \quad (3)$$

Then weak posterior consistency holds for Q on \mathcal{D}_η .

Remark 3. A similar result for jump diffusions with unit diffusion coefficient was established in [Koskela et al., 2015b, Theorem 1], and our proof will follow a similar structure. Before presenting the proof, let us highlight how the present result differs from the jump diffusion case. Both proofs of consistency require verification of a Kullback-Leibler condition for the prior, and uniform equicontinuity of the family of semigroups corresponding to densities supported by the prior. The former result is immediate by the same argument used to prove [Koskela et al., 2015b, Lemma 3], whose statement is provided below in Lemma 2 in the interest of a self-contained proof. The latter, Lemma 3 below, is different to its counterpart, [Koskela et al., 2015b, Lemma 4], which relies on positive definiteness of the diffusion coefficient.

Proof. For fixed $m \in \mathbb{N}$, the same argument used to prove Theorem 1 yields that the following convergence holds \mathbb{P}^{ϕ_0} -a.s. as $n \rightarrow \infty$:

$$\lim_{n \rightarrow \infty} Q(d\phi | \mathbf{n}_0, \dots, \mathbf{n}_m) \propto \pi^\phi(\mathbf{x}_0) \prod_{i=1}^m p_\Delta^\phi(\mathbf{x}_{i-1}, \mathbf{x}_i) Q(d\phi).$$

Hence it is sufficient to establish posterior consistency for $m + 1$ exact observations from a stationary Λ -Fleming-Viot process as $m \rightarrow \infty$. We achieve this by adapting the proof of [Koskela et al., 2015b, Theorem 1], which entails verifying two conditions. The first is that the prior places sufficient mass in Kullback-Leibler neighbourhoods of ϕ_0 , i.e. that for any $\varepsilon > 0$ we have

$$Q(\phi \in \mathcal{D}_\eta : \text{KL}(\phi_0, \phi) < \varepsilon) > 0, \quad (4)$$

where KL is Kullback-Leibler divergence between $p_{\Delta}^{\phi_0}$ and p_{Δ}^{ϕ} :

$$\text{KL}(\phi_0, \phi) := \int_{\mathcal{S}_d} \int_{\mathcal{S}_d} \log\left(\frac{p_{\Delta}^{\phi_0}(\mathbf{x}, \mathbf{y})}{p_{\Delta}^{\phi}(\mathbf{x}, \mathbf{y})}\right) p_{\Delta}^{\phi_0}(\mathbf{x}, \mathbf{y}) \pi^{\phi_0}(\mathbf{x}) d\mathbf{y} d\mathbf{x}.$$

The second is establishing uniform equicontinuity of $\{P_{\Delta}^{\phi} : \phi \in \mathcal{D}_\eta\}$:

$$\sup_{\phi \in \mathcal{D}_\eta} \sup_{\|\mathbf{x} - \mathbf{y}\|_2 < \delta} |P_{\Delta}^{\phi} f(\mathbf{x}) - P_{\Delta}^{\phi} f(\mathbf{y})| < \varepsilon$$

for each $\Delta > 0$ and $f \in \text{Lip}(\mathcal{S}_d)$, the set of Lipschitz functions on \mathcal{S}_d .

Condition (4) follows from a straightforward modification of [Koskela et al., 2015b, Lemma 3]. A statement of this result, adapted to the present context and notation, is provided below. Its proof follows the same structure that of as [Koskela et al., 2015b, Lemma 3], and is omitted.

Lemma 2. Condition (3) implies condition (4) for any $\varepsilon > 0$.

It remains to establish uniform equicontinuity on the semigroup $\{P_\Delta^\phi f : \phi \in \mathcal{D}_\eta\}$ for $f \in \text{Lip}(\mathcal{S}_d)$. This can be done by verifying the hypotheses of [Koskela et al., 2015b, Lemma 4], which we will do below.

Lemma 3. For each $\Delta > 0$ and $f \in \text{Lip}(\mathcal{S}_d)$, the collection $\{P_\Delta^\phi f : \phi \in \mathcal{D}_\eta\}$ is uniformly equicontinuous: for any $\varepsilon > 0$ there exists $\delta := \delta(\varepsilon, f, \Delta) > 0$ such that

$$\sup_{\phi \in \mathcal{D}_\eta} \sup_{\|\mathbf{x} - \mathbf{y}\|_2 < \delta} |P_\Delta^\phi f(\mathbf{x}) - P_\Delta^\phi f(\mathbf{y})| < \varepsilon.$$

Proof. The result follows by a coupling argument of [Wang, 2010, Section 3], provided the coefficients of the generator \mathcal{G}^ϕ are Lipschitz continuous for each $\phi \in \mathcal{D}_\eta$.

It will be convenient to write the generator (1) in the following form:

$$\begin{aligned} \mathcal{G}^\phi f(\mathbf{x}) &= \theta \sum_{i,j=1}^d x_j (M_{ji} - \delta_{ij}) \frac{\partial}{\partial x_i} f(\mathbf{x}) \\ &\quad + \int_\eta^1 \int_0^1 \{f(\mathbf{x} + r\Psi(u, \mathbf{x})) - f(\mathbf{x}) - r\Psi(u, \mathbf{x})\nabla f(\mathbf{x})\} r^{-2} \phi(r) du dr, \end{aligned}$$

where $\Psi(u, \mathbf{x}) := \sum_{i=1}^d (\mathbb{1}_{(S_{i-1}^\mathbf{x}, S_i^\mathbf{x}]}(u) - x_i) \mathbf{e}_i$, $S_i^\mathbf{x} := \sum_{j=1}^i x_j$, $S_0^\mathbf{x} = 0$ and \mathbf{e}_i is the canonical unit vector with a 1 in the i^{th} position and zeros elsewhere. As pointed out by Bertoin and Le Gall [2005, page 314], these generators coincide because $\int_0^1 \Psi(u, \mathbf{x}) du = 0$ for any $\mathbf{x} \in \mathcal{S}_d$.

For functions $f \in C^2(\mathcal{S}_d \times \mathcal{S}_d)$ let \mathcal{L}^ϕ be the coupling generator

$$\begin{aligned} \mathcal{L}^\phi f(\mathbf{x}, \mathbf{y}) &:= \theta \sum_{i,j=1}^d (M_{ji} - \delta_{ij}) \left(x_j \frac{\partial}{\partial x_i} f(\mathbf{x}, \mathbf{y}) + y_j \frac{\partial}{\partial y_i} f(\mathbf{x}, \mathbf{y}) \right) + \int_\eta^1 \int_0^1 r^{-2} du \phi(r) dr \\ &\quad \times \{f(\mathbf{x} + r\Psi(u, \mathbf{x}), \mathbf{y} + r\Psi(u, \mathbf{y})) - f(\mathbf{x}, \mathbf{y}) - r\Psi(u, \mathbf{x})\nabla_{\mathbf{x}} f(\mathbf{x}, \mathbf{y}) - r\Psi(u, \mathbf{y})\nabla_{\mathbf{y}} f(\mathbf{x}, \mathbf{y})\}, \end{aligned}$$

where $\nabla_{\mathbf{x}}$ and $\nabla_{\mathbf{y}}$ are the gradient operators with respect to the \mathbf{x} and \mathbf{y} variables, respectively.

Fix the test function $g(z) := z(1+z)^{-1}$. It is straightforward to check that for this function we have

$$\mathcal{L}^\phi g(\|\mathbf{x} - \mathbf{y}\|_2) \leq \frac{2B(\mathbf{x}, \mathbf{y}) + \int_0^1 \|\Psi(u, \mathbf{x}) - \Psi(u, \mathbf{y})\|_2^2 du}{2\|\mathbf{x} - \mathbf{y}\|_2^2 (1 + \|\mathbf{x} - \mathbf{y}\|_2)} g(\|\mathbf{x} - \mathbf{y}\|_2) =: K(\mathbf{x}, \mathbf{y}) g(\|\mathbf{x} - \mathbf{y}\|_2)$$

uniformly in $\phi \in \mathcal{D}_\eta$, where $B(\mathbf{x}, \mathbf{y}) := \theta(\mathbf{x} - \mathbf{y})^T (M - \mathbb{I}_d)(\mathbf{x} - \mathbf{y})$ and \mathbb{I}_d is the $d \times d$ identity matrix.

Proposition 2.2 of [Wang, 2010] requires that $K(\mathbf{x}, \mathbf{y})\|\mathbf{x} - \mathbf{y}\|_2 \leq C_\delta \|\mathbf{x} - \mathbf{y}\|_2$ for some constant C_δ depending only on δ , θ and M . Note that $B(\mathbf{x}, \mathbf{y}) \leq \theta \|M - \mathbb{I}_d\|_2 \|\mathbf{x} - \mathbf{y}\|_2^2$, so that the result will follow if $\int_0^1 \|\Psi(u, \mathbf{x}) - \Psi(u, \mathbf{y})\|_2^2 du \leq C \|\mathbf{x} - \mathbf{y}\|_2$ for some constant C . Now

$$\begin{aligned} \int_0^1 \|\Psi(u, \mathbf{x}) - \Psi(u, \mathbf{y})\|_2^2 du &= 2 - \|\mathbf{x} - \mathbf{y}\|_2^2 - 2 \sum_{i=1}^d \int_0^1 \mathbb{1}_{(S_{i-1}^\mathbf{x}, S_i^\mathbf{x}]}(u) \mathbb{1}_{(S_{i-1}^\mathbf{y}, S_i^\mathbf{y}]}(u) du \\ &= 2 \sum_{i=1}^d \left(S_i^\mathbf{x} - S_{i-1}^\mathbf{x} - \int_0^1 \mathbb{1}_{(S_{i-1}^\mathbf{x}, S_i^\mathbf{x}]}(u) \mathbb{1}_{(S_{i-1}^\mathbf{y}, S_i^\mathbf{y}]}(u) du \right) - \|\mathbf{x} - \mathbf{y}\|_2^2, \end{aligned}$$

and hence it suffices to bound $\sum_{i=1}^d S_i^\mathbf{x} - S_{i-1}^\mathbf{x} - \int_0^1 \mathbb{1}_{(S_{i-1}^\mathbf{x}, S_i^\mathbf{x}]}(u) \mathbb{1}_{(S_{i-1}^\mathbf{y}, S_i^\mathbf{y}]}(u) du \leq C \|\mathbf{x} - \mathbf{y}\|_2$. Consider a single index $i \in [d]$ and suppose without loss of generality that $S_i^\mathbf{y} \leq S_i^\mathbf{x}$ and if

$S_i^y = S_i^x$ then $S_{i-1}^x < S_{i-1}^y$. There are two cases: either $(S_{i-1}^y, S_i^y] \subseteq (S_{i-1}^x, S_i^x]$ or $S_{i-1}^y < S_{i-1}^x$ and $S_i^y < S_i^x$. In the former we have

$$S_i^x - S_{i-1}^x - \int_0^1 \mathbb{1}_{(S_{i-1}^x, S_i^x]}(u) \mathbb{1}_{(S_{i-1}^y, S_i^y]}(u) du = x_i - y_i \leq \|\mathbf{x} - \mathbf{y}\|_1,$$

while in the latter we have

$$\begin{aligned} S_i^x - S_{i-1}^x - \int_0^1 \mathbb{1}_{(S_{i-1}^x, S_i^x]}(u) \mathbb{1}_{(S_{i-1}^y, S_i^y]}(u) du &= S_i^x - S_{i-1}^x - (S_i^y - S_{i-1}^y)^+ \\ &\leq S_i^x - S_i^y + S_{i-1}^x - S_{i-1}^y \leq \sum_{j=i+1}^d |x_j - y_j| + \sum_{j=1}^{i-1} |x_j - y_j| \leq \|\mathbf{x} - \mathbf{y}\|_1 \end{aligned}$$

because the L^1 -distance between the sub-probability measures (x_1, \dots, x_{i-1}) and (y_1, \dots, y_{i-1}) on $[i-1]$ is at least as large as the difference in their total masses, and likewise for the corresponding pair on $\{i+1, \dots, d\}$. Summing over indices yields

$$\int_0^1 \|\Psi(u, \mathbf{x}) - \Psi(u, \mathbf{y})\|_2^2 du \leq 2d^{3/2} \|\mathbf{x} - \mathbf{y}\|_2 \quad (5)$$

since $\|\mathbf{z}\|_1 \leq \sqrt{d} \|\mathbf{z}\|_2$ by the Cauchy-Schwarz inequality. Hence by [Wang, 2010, Proposition 2.2] we have that

$$\sup_{\|\mathbf{x} - \mathbf{y}\|_2 < \delta} |P_\Delta^\phi f(\mathbf{x}) - P_\Delta^\phi f(\mathbf{y})| < \varepsilon.$$

Uniformity in $\phi \in \mathcal{D}_\eta$ follows because the bound on the R.H.S. of (5) does not depend on ϕ . \square

The remainder of the proof follows as in [Koskela et al., 2015b]. It suffices to show that for $f \in \text{Lip}(\mathcal{S}_d)$ and $B := \{\phi \in \mathcal{D}_\eta : \|P_\Delta^\phi f - P_\Delta^{\phi_0} f\|_{1,\nu} > \varepsilon\}$ we have $Q(B|\mathbf{x}_0, \dots, \mathbf{x}_m) \rightarrow 0$ \mathbb{P}^{ϕ_0} -almost surely. To that end we fix $f \in \text{Lip}(\mathcal{S}_d)$, $\varepsilon > 0$ and thus the set B . Condition (4) implies that [van der Meulen and van Zanten, 2013, Lemma 5.2] holds, so that if for a measurable collection of subsets $C_m \subset \mathcal{D}_\eta$ there exists $c > 0$ such that

$$e^{mc} \int_{C_m} \pi^\phi(\mathbf{x}_0) \prod_{i=1}^m p_\Delta^\phi(\mathbf{x}_{i-1}, \mathbf{x}_i) Q(d\phi) \rightarrow 0$$

\mathbb{P}^{ϕ_0} -almost surely, then $Q(C_m|\mathbf{x}_0, \dots, \mathbf{x}_m) \rightarrow 0$ \mathbb{P}^{ϕ_0} -almost surely as well. Likewise, Lemma 3 implies [van der Meulen and van Zanten, 2013, Lemma 5.3]: there exists a compact subset $K \subset \mathcal{S}_d$, $N \in \mathbb{N}$ and compact, connected sets I_1, \dots, I_N that cover K such that

$$B \subset \bigcup_{j=1}^N B_j^+ \cup \bigcup_{j=1}^N B_j^-,$$

where

$$B_j^\pm := \left\{ \phi \in \mathcal{D}_\eta : P_\Delta^\phi f(\mathbf{x}) - P_\Delta^{\phi_0} f(\mathbf{x}) > \pm \frac{\varepsilon}{4\nu(K)} \text{ for every } \mathbf{x} \in I_j \right\}.$$

Thus it is only necessary to show $Q(B_j^\pm|\mathbf{x}_0, \dots, \mathbf{x}_m) \rightarrow 0$ \mathbb{P}^{ϕ_0} -almost surely. Define the stochastic process

$$D_m := \left(\int_{B_j^+} \pi^\phi(\mathbf{x}_0) \prod_{i=1}^m p_\Delta^\phi(\mathbf{x}_{i-1}, \mathbf{x}_i) Q(d\phi) \right)^{1/2}.$$

Now $D_m \rightarrow 0$ exponentially fast as $m \rightarrow \infty$ by an argument identical to that used to prove [van der Meulen and van Zanten, 2013, Theorem 3.5]. The same is also true of the analogous stochastic process defined by integrating over B_j^- , which completes the proof. \square

Note that bounding the support of Λ away from 0 is not necessary: the proof could be adapted to equivalent collections of measures supported on $(0, 1]$ for which (3) holds, e.g. by requiring $\phi(r) \sim r^2$ as $r \searrow 0$ for Q -a.e. $\phi \in \mathcal{D}_\eta$. However, any such condition is bound to be mathematically restrictive. In the next section we show that given a data set of size n , it is natural to infer the first $n - 2$ moments of Λ because they fully capture the signal in the data set. Example 2 at the end of Section 4 provides a family of priors which satisfy the hypotheses of Theorem 2, and whose support can be chosen to contain arbitrarily close approximations to truncated moment sequences of any $\Lambda \in \mathcal{M}_1([0, 1])$.

Remark 4. The hypotheses of Theorem 2 are strong, and thus it would be desirable to obtain a posterior contraction rate in addition to just consistency. In fact, methods akin to that employed in the proof have been extended to provide rates for compound Poisson processes [Gugushvili et al., 2015] and scalar diffusions on compact intervals [Nickl and Söhl, 2015]. However, extending either approach to our setting would require bounds of the form

$$\|\pi^\phi - \pi^{\phi_0}\|_2 \leq Cn^{-\beta}, \quad \|\pi^\phi / \pi^{\phi_0} - 1\|_2 \leq \tilde{C}n^{-\tilde{\beta}}$$

for some constants $\beta, \tilde{\beta}, C, \tilde{C} > 0$. Since the Λ -Fleming-Viot stationary density is intractable in nearly all cases, it does not seem possible to extend our approach to obtain rates of posterior consistency.

4 Parametrisation by finitely many moments

Consider a set of type frequencies $\mathbf{n} \in \mathbb{N}^d$ of size $n := \sum_{i=1}^d n_i$ generated by a Λ -coalescent with finite alleles mutation started from ψ_n .

Lemma 4. *The likelihood satisfies $\mathbf{P}_n^\Lambda(\mathbf{n}) = \mathbf{P}_n^{\lambda_3, \lambda_4, \dots, \lambda_n}(\mathbf{n})$. That is, Λ is conditionally independent of \mathbf{n} given $\{\lambda_k\}_{k=3}^n$.*

Proof. Let $-q_{nn} = \sum_{k=1}^{n-1} \binom{n}{n-k+1} \lambda_{n, n-k+1}$ be the total merger rate of the Λ -coalescent with n blocks. It is well known that the Λ -coalescent likelihood is the unique solution to the recursion [Möhle, 2006, Birkner and Blath, 2008, 2009]:

$$\begin{aligned} \mathbf{P}_n^\Lambda(\mathbf{n}) &= \frac{\theta}{n\theta - q_{nn}} \sum_{i,j=1}^d (n_j - 1 + \delta_{ij}) M_{ji} \mathbf{P}_n^\Lambda(\mathbf{n} - \mathbf{e}_i + \mathbf{e}_j) \\ &\quad + \frac{1}{n\theta - q_{nn}} \sum_{i: n_i \geq 2} \sum_{k=2}^{n_i} \binom{n_i}{k} \lambda_{n_i, k} \frac{n_i - k + 1}{n - k + 1} \mathbf{P}_{n-k+1}^\Lambda(\mathbf{n} - (k-1)\mathbf{e}_i), \end{aligned} \quad (6)$$

with boundary condition $\mathbf{P}_1^\Lambda(\mathbf{e}_i) = m(i)$. Repeated application of this recursion yields a closed system of linear equations for the likelihood because all sample sizes on the R.H.S. are equal to or smaller than the one on the left hand side. This system is far too large to solve for all but very small sample sizes, but it is clear that the solution can depend on Λ only through the polynomial moments $\{\lambda_{m,k}\}_{k \leq m=2}^n$.

Polynomial moments can be written as a linear combination of monomial moments:

$$\lambda_{m,k} = \sum_{j=0}^{m-k} \binom{m-k}{j} (-1)^j \lambda_{k+j}, \quad (7)$$

which means that only the monomial moments $\{\lambda_k\}_{k=2}^n$ are required. Since $\lambda_2 = \int_0^1 \Lambda(dx) = 1$, the moments $\{\lambda_k\}_{k=3}^n$ are sufficient. \square

Motivated by Lemma 4 we make the following definition:

Definition 4. Let \sim_n be the equivalence relation on $\mathcal{M}_1([0, 1])$ defined via

$$\Lambda_1 \sim_n \Lambda_2 \text{ if } \lambda_k^{(1)} = \lambda_k^{(2)} \text{ for } k \in \{3, \dots, n\}$$

where $\lambda_k^{(i)} := \int_0^1 x^{k-2} \Lambda_i(dx)$. We call the equivalence classes of \sim_n moment classes of order n .

In view of Lemma 4 it is natural to consider the problem of inferring Λ from \mathbf{n} in the quotient space $\mathcal{M}_1([0, 1]) / \sim_n$, not in $\mathcal{M}_1([0, 1])$. Moreover, requiring all linear combinations of the form (7) to be non-negative guarantees a unique solution to the Hausdorff moment problem, so that each completely monotonic moment sequence bounded by 1 corresponds to some $\Lambda \in \mathcal{M}_1([0, 1])$. Hence we parametrise the space $\mathcal{M}_1([0, 1]) / \sim_n$ by truncated, completely monotonic moment sequences of length $n - 2$ with leading term $\lambda_3 \leq 1$. This approach yields a compact, finite-dimensional parameter space which nevertheless captures all the signal in the data. Table 2 lists some moment sequences corresponding to popular families of Λ -measures.

Λ	δ_0	δ_1	Beta($2 - \alpha, \alpha$)	$U(0, 1)$	$\frac{2}{2+\psi^2}\delta_0 + \frac{\psi^2}{2+\psi^2}\delta_\psi$	$c\delta_0 + \frac{(1-c)}{2}rdr$
λ_k	0	1	$\frac{(2-\alpha)_{k-2}}{(2)_{k-2}}$	$\frac{1}{k-1}$	$\frac{\psi^k}{2+\psi^2}$	$\frac{1-c}{2k}$

Table 2: Moment sequences of particular Λ -coalecscents. Here $(a)_k := a(a+1)\dots(a+k-1)$ denotes the rising factorial.

Naturally, the prior Q ought to be chosen to yield a tractable push-forward prior on truncated moment sequences. These push-forward priors inherit posterior consistency whenever Q satisfies the conditions of Theorem 2 because truncated moment sequences can be written as bounded functionals of Λ .

5 Prior distributions

In this section we provide an example family of priors which satisfies the consistency criteria of Theorem 2 and has tractable push-forward distributions on truncated moment sequences.

Definition 5. Let $Q \in \mathcal{M}_1(\mathcal{M}_1([0, 1]))$ be a prior distribution for Λ . Then the moments $\{\lambda_k\}_{k=3}^n$ have joint prior Q_n on the space of completely monotonic sequences of length $n - 2$ given by

$$Q_n(\lambda_3 \in dy_3, \dots, \lambda_n \in dy_n) := \int_{\mathcal{M}_1([0, 1])} \prod_{k=3}^n \mathbb{1}_{\{dy_k\}} \left(\int_{(0, 1]} r^{k-2} \Lambda(dr) \right) Q(d\Lambda). \quad (8)$$

The prior Q should to be chosen such that the R.H.S. of (8) is tractable, and the following example illustrates that such a choice is possible.

Example 2. Fix $\eta > 0$ and $\alpha \in \mathcal{M}([\eta, 1])$ with finite mass and a strictly positive Lebesgue density $\alpha(r)$. Suppose \mathcal{D}_η satisfies the conditions of Definition 1, and in addition that every $\phi \in \mathcal{D}_\eta$ is continuous. Let $R(d\tau)$ be a probability measure on $(0, \infty)$ placing positive mass in all non-empty open sets. For $x \in [\eta, 1]$ and $\tau > 0$ let

$$q_{x, \tau}(r) := \frac{\mathbb{1}_{[\eta, 1]}(r - x) h_\tau(r - x)}{h_\tau([\eta, 1])},$$

where h_τ is the Gaussian density on \mathbb{R} with mean 0 and variance τ^{-1} .

Let $DP(\alpha)$ be the law of a Dirichlet process centred on α [Ferguson, 1973] and let Q be given by the Dirichlet process mixture distribution [Lo, 1984] with mixing distribution $DP(\alpha) \otimes R$ and mixture components $q_{x,\tau}$. In other words, let $F \sim DP(\alpha)$ and $F(r_1), F(r_2), \dots$ be the weights of ordered atoms of F . Let $\{\tau_i\}_{i=1}^\infty$ be i.i.d. draws from R . Then a draw from Q is given by

$$\phi(r)|F, \{\tau_i\}_{i=1}^\infty = \sum_{i=1}^\infty F(r_i) q_{r_i, \tau_i}(r).$$

The prior Q places full mass on equivalent densities bounded from above and away from 0 by construction. It remains to check (3) to verify Q has posterior consistency. Theorem 1 of [Bhattacharya and Dunson, 2012] yields that for any $\delta > 0$ the prior places positive mass in all open balls: $Q(\phi \in \mathcal{D}_\eta : \|\phi - \phi_0\|_\infty < \delta) > 0$ for any $\phi_0 \in \mathcal{D}_\eta$ under the above assumptions on α , R and $(q_{x,\tau})_{x \in [\eta, 1], \tau > 0}$. Now fix $\varepsilon > 0$, $\phi_0 \in \mathcal{D}_\eta$, as well as $0 < \delta < c =: \inf_{r \in [\eta, 1]} \{\phi_0(r)\}$ such that

$$\log\left(\frac{c}{c-\delta}\right) \vee \left|\log\left(\frac{c}{c+\delta}\right)\right| + \frac{\delta}{c-\delta} < \eta^2 \varepsilon.$$

Then

$$\begin{aligned} Q\left(\phi \in \mathcal{D}_\eta : \int_\eta^1 \left\{ \left| \log\left(\frac{\phi_0(r)}{\phi(r)}\right) \right| + \left| \frac{\phi_0(r)}{\phi(r)} - 1 \right| \right\} r^{-2} \phi_0(r) dr < \varepsilon\right) \\ \geq Q(\phi \in \mathcal{D}_\eta : \|\phi - \phi_0\|_\infty < \delta) > 0, \end{aligned}$$

because for any ϕ satisfying $\|\phi - \phi_0\|_\infty < \delta$ we have

$$\begin{aligned} & \int_\eta^1 \left\{ \left| \log\left(\frac{\phi_0(r)}{\phi(r)}\right) \right| + \left| \frac{\phi_0(r)}{\phi(r)} - 1 \right| \right\} r^{-2} \phi_0(r) dr \\ & \leq \int_\eta^1 \left\{ \log\left(\frac{\phi_0(r)}{\phi_0(r)-\delta}\right) \vee \left| \log\left(\frac{\phi_0(r)}{\phi_0(r)+\delta}\right) \right| + \left| \frac{\phi_0(r)}{\phi_0(r)-\delta} - 1 \right| \right\} r^{-2} \phi_0(r) dr \\ & \leq \left(\log\left(\frac{c}{c-\delta}\right) \vee \left| \log\left(\frac{c}{c+\delta}\right) \right| + \frac{\delta}{c-\delta} \right) \eta^{-2} < \varepsilon. \end{aligned}$$

We now use the machinery of Regazzini et al. [2002] to give an explicit system of equations for the distribution function of Q_n under this choice of Q . Define the family of functions

$$g_p(x) := \int_\eta^1 r^p q_{x,\tau}(r) dr$$

for $p \in \mathbb{N}$ and $x \in [\eta, 1]$, as well as the vectors $\mathbf{g}_n(x) := (g_1(x), g_2(x), \dots, g_n(x))$ and $\mathbf{s}_n := (s_1, \dots, s_n) \in \mathbb{R}^n$. For brevity, for a measure ν and a function f let $\nu(f) := \int f d\nu$ whenever the integral exists.

Let γ_α be a Gamma random measure with parameter α , that is, a random finite measure on $[\eta, 1]$ such that for any measurable partition $\{A_1, \dots, A_n\}$ the random variables $(\gamma_\alpha(A_1), \dots, \gamma_\alpha(A_n))$ are independent and gamma distributed with common scale parameter 1 and respective shape parameters $\alpha(A_k)$. Let

$$h_n(\mathbf{s}_n; \mathbf{g}_n; \alpha) := \mathbb{E}[\exp(i\mathbf{s}_n \cdot \gamma_\alpha(\mathbf{g}_n))]$$

be the characteristic function of $\gamma_\alpha(\mathbf{g}_n) := (\gamma_\alpha(g_1), \dots, \gamma_\alpha(g_n))$. Note that $h_n(\mathbf{s}_n; \mathbf{g}_n; \alpha) = h_n(\mathbf{1}; \mathbf{s}_n \cdot \mathbf{g}_n; \alpha)$ and by [Regazzini et al., 2002, Proposition 10]

$$h_n(\mathbf{s}_n; \mathbf{g}_n; \alpha) = \exp\left(-\int_0^1 \log(1 - i\mathbf{s}_n \cdot \mathbf{g}_n) d\alpha\right). \quad (9)$$

Now let $F_n(\boldsymbol{\sigma}, \mathbf{g}_n, \alpha)$ be the joint distribution function of $F(\mathbf{g}_n) := (F(g_1), \dots, F(g_n))$ under $DP(\alpha)$. The following trick was introduced in [Hannum et al., 1981, equation (2.9)]:

$$F_n(\boldsymbol{\sigma}, \mathbf{g}_n, \alpha) = F_n(\mathbf{0}, \gamma_\alpha(\mathbf{g}_n - \boldsymbol{\sigma}), \alpha)$$

for any $\boldsymbol{\sigma} \in \mathbb{R}^n$, so that it is sufficient to invert h_n at the origin to obtain F_n . This can be done using the multidimensional version of the Gurland inversion formula [Gurland, 1948, Theorem 3]:

Let $C_0, \dots, C_n \in \mathbb{R}^{n+1}$ solve

$$\begin{aligned} C_n &= -1 \\ \sum_{k=0}^{n-r-1} \binom{n-r}{k} C_{r+k} &= 1 \text{ for } r \in \{0, \dots, n-1\}. \end{aligned} \quad (10)$$

Then

$$\begin{aligned} (-1)^{n+1} 2^n F_n(\boldsymbol{\sigma}, \mathbf{g}_n, \alpha) &= C_0 \\ &+ \sum_{k=1}^n \frac{C_k}{(\pi i)^k} \sum_{1 \leq j_1 < \dots < j_k \leq n} \int_0^\infty \dots \int_0^\infty \frac{h_k(\mathbf{s}_k; g_{j_1} - \sigma_{j_1}, \dots, g_{j_k} - \sigma_{j_k}; \alpha)}{s_1 \times \dots \times s_k} d\mathbf{s}_k. \end{aligned}$$

The characteristic functions h_k and the constants C_k can be computed from (9) and (10) respectively, so that the R.H.S. can be evaluated numerically for practical applications. Numerical methods are discussed in [Regazzini et al., 2002, Section 6].

Finally, we demonstrate that the restrictive assumptions of Theorem 2 still allow inference for broad classes of moment sequences with arbitrarily small approximation errors. Let $\beta(r)$ be any non-negative probability density on $[0, 1]$, and define the truncation $\bar{\beta}(r) := \kappa(\eta, c, K)^{-1}(\beta(r) \vee c \wedge K)$, where κ is the normalising constant

$$\kappa(\eta, c, K) := \int_\eta^1 \beta(r) - (\beta(r) - K)^+ + (c - \beta(r))^+ dr.$$

Note that $Q(\phi \in \mathcal{D}_\eta : \|\phi - \bar{\beta}\|_\infty < \delta) > 0$ for any $\delta > 0$, and fix such a ϕ . Now consider the error on the k^{th} moment:

$$\begin{aligned} \left| \int_0^1 r^k \beta(r) dr - \int_\eta^1 r^k \phi(r) dr \right| &\leq \\ \beta((0, \eta)) + (1 - \eta)\delta &+ \frac{(\kappa(\eta, c, K) - 1)\beta([\eta, 1]) + c(1 - \eta)}{\kappa(\eta, c, K)} + \int_\eta^1 \frac{(\beta(r) - K)^+}{\kappa(\eta, c, K)} dr. \end{aligned}$$

Each term on the R.H.S. can be made small by choosing η , c and δ sufficiently small, and K sufficiently large because $\kappa \rightarrow 1$ as $\eta \rightarrow 0$, $c \rightarrow 0$ and $K \rightarrow \infty$, and

$$\int_\eta^1 (\beta(r) - K)^+ dr = 1 - \int_\eta^1 \beta(r) \wedge K dr \rightarrow \beta((0, \eta))$$

as $K \rightarrow \infty$ by the Monotone Convergence Theorem. A further approximation step also enables consideration of atoms by choosing β which places all of its mass in neighbourhoods of the desired locations for atoms. Hence it is possible to ensure the support of Q extends arbitrarily close to any desired moment sequences despite the restrictive assumptions on \mathcal{D}_η in Theorem 2.

6 Robust bounds on functionals of Λ

Having established consistency criteria for the posterior and a finite parametrisation via $n - 2$ leading moments, we now turn to what can be said about Λ based on inferring these moments. It would be ideal if the diameter of moment classes shrunk with increasing n , as then it would be possible to fix a representative $\Lambda \in \mathcal{M}_1([0, 1])$ with specified $n - 2$ leading moments and control the remaining within-moment-class error. In Theorem 3 we show that such shrinking does not happen, and devote the remainder of the section to presenting quantities which can be controlled based on $n - 2$ moments alone. We begin by recalling some standard results from the theory of orthogonal polynomials.

Definition 6. Suppose n is odd. Let $m := \frac{n-3}{2}$ and $\{\phi_k\}_{k=0}^m$ be the first $m + 1$ Λ -orthogonal polynomials. Let $\{\xi_k\}_{k=1}^m$ be the zeros of ϕ_m .

Remark 5. It is a standard result that $\{\phi_k\}_{k=0}^{m-1}$ and $\{\xi_k\}_{k=1}^m$ are constant within moment classes of order at least n .

The following bounds on Λ in terms of its leading $n - 2$ moments are classical:

Lemma 5 (Chebyshev-Markov-Stieltjes (CMS) inequalities). *Define*

$$\rho_{m-1}(z) := \left(\sum_{k=0}^{m-1} |\phi_k(z)|^2 \right)^{-1}.$$

Then the following inequalities are sharp:

$$\Lambda([0, \xi_j]) \leq \sum_{k=1}^j \rho_{m-1}(\xi_k) \leq \Lambda([0, \xi_{j+1}]) \text{ for } j = \{1, \dots, m\},$$

where $\xi_{m+1} := 1$.

We are now in a position to prove the following theorem:

Theorem 3. *For any $n \in \mathbb{N}$ and any completely monotonic sequence of moments $\{\lambda_{k,k}\}_{k=3}^n$ with $\lambda_{3,3} \leq 1$ there exist uncountably many measures $\Lambda \in \mathcal{M}_1([0, 1])$, all with leading moments $\{\lambda_{k,k}\}_{k=3}^n$ and all satisfying the CMS inequalities, such that for any pair Λ_x and Λ_y of them we have $d_{TV}(\Lambda_x, \Lambda_y) = 2$, where d_{TV} is the total variation distance.*

Proof. It will be convenient to write the CMS inequalities in the following, equivalent form:

$$\begin{aligned} 0 &\leq \Lambda([0, \xi_1]) \leq \rho_{m-1}(\xi_1) \\ 0 &\leq \Lambda([\xi_j, \xi_{j+1}]) \leq \rho_{m-1}(\xi_j) + \rho_{m-1}(\xi_{j+1}) \text{ for } j \in \{1, \dots, m-1\} \\ 0 &\leq \Lambda([\xi_m, 1]) \leq \rho_{m-1}(\xi_m) \end{aligned}$$

where the last inequality follows from the fact that $\sum_{k=1}^m \rho_{m-1}(\xi_k) = 1$. This equality holds because $\sum_{k=1}^m \rho_{m-1}(\xi_k)$ is the sum of all order m Gauss quadrature weights, or equivalently the quadrature applied to the constant function 1, which is a polynomial of degree 0. The equality follows by recalling that Gauss quadrature is exact for polynomials of order up to $2m - 1$.

Now let the measures Λ_x and Λ_y be described by sequences of $(m + 1)$ weights (x_0, x_1, \dots, x_m) and (y_0, y_1, \dots, y_m) , with the j^{th} weight denoting the mass that the corresponding measure places in the interval $[\xi_j, \xi_{j+1}]$ (with obvious adjustments for the right hand boundary terms).

For brevity let $\zeta_j := \rho_{m-1}(\xi_j)$. Suppose first that m is odd, and let the vectors of weights be given as

$$\begin{aligned}(x_0, x_1, x_2, x_3, x_4, \dots, x_{m-1}, x_m) &= (\zeta_1, 0, \zeta_2 + \zeta_3, 0, \zeta_4 + \zeta_5, \dots, 0, \zeta_m) \\ (y_0, y_1, y_2, y_3, y_4, \dots, y_{m-1}, y_m) &= (0, \zeta_1 + \zeta_2, 0, \zeta_3 + \zeta_4, 0, \dots, \zeta_{m-1} + \zeta_m, 0)\end{aligned}$$

Both measures have total mass $\sum_{j=1}^m \zeta_j = 1$, and the interlacing masses have no overlap so $d_{TV}(\Lambda_x, \Lambda_y) = 2$. The case where m is even is similar. \square

Remark 6. The same result holds in Kullback-Leibler divergence due to Pinsker's inequality: for probability measures P and Q such that $P \sim Q$, and letting $\text{KL}(P, Q) := \mathbb{E}_P[\log(dP/dQ)]$ we have

$$d_{TV}(P, Q) \leq \sqrt{\frac{1}{2} \text{KL}(P, Q)},$$

so that $\text{KL}(\Lambda_x, \Lambda_y) \geq 8$ for Λ_x and Λ_y as in Theorem 3.

Despite this seemingly disappointing result, it is possible to make some conclusions about Λ based on $n - 2$ moments. For example, the Kingman hypothesis can be tested in a robust way by checking whether the vector $(0, 0, \dots, 0)$ lies in a desired credible region of the posterior $Q_n(\cdot | \mathbf{n})$, and the plausibility of any other Λ of interest can be assessed similarly. More generally, it is possible to maximise/minimise a certain class of functionals subject to moment constraints obtained from a credible region to obtain robust bounds for quantities of interest. We begin by recalling some relevant definitions.

Definition 7. Let $m, n \in \mathbb{N}$, and fix \mathbb{R} -valued constants $\{c_k\}_{k=1}^m$, a sequence $\{i_k\}_{k=1}^m$ of $\{3, \dots, n\}$ -valued indices and a binary sequence $\{j_k\}_{k=1}^m$ of zeros and ones. Let

$$\mathcal{C}_m := \{\Lambda \in \mathcal{M}_1([0, 1]) : (-1)^{j_k} \lambda_{i_k} \leq c_k \text{ for } k \in \{1, \dots, m\}\} \quad (11)$$

be a subset of $\mathcal{M}_1([0, 1])$ with leading $n - 2$ moments in a desired region specified by m linear inequalities. Let $\text{ext}(\mathcal{C}_m)$ be the extremal points in \mathcal{C}_m , i.e. those which cannot be written as non-trivial convex combinations of elements in \mathcal{C}_m , and

$$\mathcal{C}_D := \left\{ \nu \in \mathcal{C}_m : \nu = \sum_{k=1}^p w_k \delta_{x_k} \text{ where } 1 \leq p \leq m + 1, w_k \geq 0, x_k \in [0, 1] \text{ and } \sum_{k=1}^p w_k = 1 \right\}$$

be the set of discrete probability measures on $[0, 1]$ with at most $m + 1$ atoms.

Example 3. The extremal points of $\mathcal{M}_1([0, 1])$ are the Dirac measures:

$$\text{ext}(\mathcal{M}_1([0, 1])) = \{\delta_x : x \in [0, 1]\}.$$

For our purposes \mathcal{C}_m should be thought of as an envelope containing a desired credible region of truncated, completely monotonic moment sequences expressed using finitely many linear constraints. We postpone discussion of how an approximate credible region can be obtained to the next section, and simply assume one is available. The importance of Definition 7 is that maxima and minima of certain functionals of Λ coincide on \mathcal{C}_m and \mathcal{C}_D , and that \mathcal{C}_D is finite-dimensional so that these extrema can be found numerically. The class of functionals for which this can be done is given below.

Definition 8. The functional $F : \mathcal{C} \mapsto \overline{\mathbb{R}}$ is *measure-affine* if, for every $\nu \in \mathcal{C}$ and $p \in \mathcal{M}_1(\text{ext}(\mathcal{C}))$ such that $\nu(E) = \int_{\text{ext}(\mathcal{C})} \gamma(E) p(d\gamma)$ for every $E \in \mathcal{B}([0, 1])$, F is p -integrable and

$$F(\nu) = \int_{\text{ext}(\mathcal{C})} F(\gamma) p(d\gamma).$$

Intuitively, ν is a barycentre of \mathcal{C} with weights on extremal points given by p , and F is measure-affine if it commutes with the operation of expressing ν as the weighted sum of extremal points. If \mathcal{C} consists of finitely many points, this definition coincides with the usual definition of affine functions.

The following two results originate from [Winkler, 1988, Proposition 3.1 and Theorem 3.2]:

Lemma 6. *If $q : [0, 1] \mapsto \overline{\mathbb{R}}$ is bounded on at least one side then $F : \nu \mapsto \mathbb{E}_\nu[q]$ is measure-affine.*

Lemma 7. *Let \mathcal{C}_m be as in Definition 7 and $F : \mathcal{C}_m \mapsto \overline{\mathbb{R}}$ be measure-affine. Then*

$$\inf_{\nu \in \mathcal{C}_m} F(\nu) = \inf_{\nu \in \mathcal{C}_D} F(\nu) \quad (12)$$

$$\sup_{\nu \in \mathcal{C}_m} F(\nu) = \sup_{\nu \in \mathcal{C}_D} F(\nu). \quad (13)$$

The purpose of Lemma 7 is that the optimisation problems on the R.H.S. of (12) and (13) are finite-dimensional and can be solved numerically. Hence tight bounds for measure-affine functionals $F(\Lambda)$ over credibility regions can be computed in an assumption-free manner.

In order to specify \mathcal{C}_m it remains to be able to approximate the posterior, which we achieve via MCMC. This will be detailed in the next section. Before that, we conclude this section with a simple example computation.

Example 4. Suppose a posterior credible region is specified via two linear constraints as

$$\mathcal{C}_2 = \{\Lambda \in \mathcal{M}_1([0, 1]) : \lambda_3 \leq 0.5 \text{ and } 0.3 \leq \lambda_4\},$$

and that the measure-affine functional of interest is the exponential:

$$F(\Lambda) := \int_0^1 e^{-r} \Lambda(dr).$$

Then the finite dimensional subspace $\mathcal{C}_D \subset \mathcal{C}_2$ consists of discrete probability measures on $[0, 1]$ with at most three atoms:

$$\mathcal{C}_D = \left\{ \Lambda \in \mathcal{C}_2 : \Lambda = \sum_{k=1}^p w_k \delta_{r_k} \text{ where } 1 \leq p \leq 3, w_k \geq 0, r_k \in [0, 1] \text{ and } \sum_{k=1}^p w_k = 1 \right\}.$$

This yields three maximisation/minimisation problems, one corresponding to each number of atoms, though in practice only the largest needs to be solved since the two others can be recovered as special cases. In this case, the constrained optimisation problem is

$$\begin{aligned} &\text{Maximise/Minimise: } ae^{-x} + be^{-y} + (1-a-b)e^{-z} \\ &\text{Subject to: } ax + by + (1-a-b)z \leq 0.5 \\ &\quad -ax^2 - by^2 - (1-a-b)z^2 \leq -0.3 \\ &\quad a, b \leq 1 \\ &\quad -a, -b \leq 0 \\ &\quad a + b \leq 1 \\ &\quad x, y, z \leq 1 \\ &\quad -x, -y, -z \leq 0. \end{aligned}$$

Numerical evaluation in Mathematica yields the bounds $F(\Lambda) \in (0.620, 0.810)$.

7 Simulation study

Efficient methods for approximating the Λ -coalescent likelihood pointwise exist [Birkner et al., 2011, Koskela et al., 2015a] and can be readily adapted to the form developed by Beaumont [2003] for time series data. These likelihood estimators can then be used in the pseudo-marginal Metropolis-Hastings algorithm [Beaumont, 2003, Andrieu and Roberts, 2009], in which the likelihood evaluations required in a standard Metropolis-Hastings algorithm are replaced with unbiased estimators. The resulting algorithm still targets the correct posterior and inherits the efficient exploration of parameter space of MCMC methods. Thus it is well-suited to high-dimensional situations with intractable likelihood.

Let $S(n)$ denote the space of completely monotonic sequences of length $n - 2$, and for $\lambda \in S(n)$ let $L(\lambda; \mathbf{n})$ be the likelihood function and $\hat{L}(\lambda; \mathbf{n})$ be an unbiased estimator. We recall the pseudo-marginal Metropolis-Hastings algorithm in Algorithm 1 below.

Algorithm 1 Pseudo-marginal Metropolis-Hastings for finite moment sequences

Require: Prior P_n , observation \mathbf{n} , transition kernel $K : S(n) \times S(n) \mapsto \mathbb{R}_+$ and $N \in \mathbb{N}$.

- 1: Initialise sample $S \leftarrow \emptyset$ and moment sequence $\lambda \leftarrow \lambda_0$.
 - 2: Compute likelihood estimator $\hat{L}(\lambda; \mathbf{n})$.
 - 3: **for** $j = 1, \dots, N$ **do**
 - 4: Sample $\lambda' \sim K(\lambda, \cdot)$.
 - 5: Compute likelihood estimator $\hat{L}(\lambda'; \mathbf{n})$.
 - 6: Set $a \leftarrow 1 \wedge \frac{K(\lambda', \lambda) \hat{L}(\lambda'; \mathbf{n}) P_n(\lambda')}{K(\lambda, \lambda') \hat{L}(\lambda; \mathbf{n}) P_n(\lambda)}$.
 - 7: Sample $U \sim U(0, 1)$.
 - 8: **if** $U < a$ **then**
 - 9: Set $S \leftarrow S \cup \{\lambda'\}$, $\hat{L}(\lambda) \leftarrow \hat{L}(\lambda')$ and $\lambda \leftarrow \lambda'$.
 - 10: **else**
 - 11: Set $S \leftarrow S \cup \{\lambda\}$.
 - 12: **end if**
 - 13: **end for**
 - 14: **return** S
-

Algorithm 1 returns a sample of moment sequences S , whose limiting distribution is the posterior. A credible region C can be approximated from MCMC output, and used to form \mathcal{C}_m as per (11). Measure-affine quantities of interest can then be maximised or minimised using finite computation by making use of Lemma 7.

By way of demonstration we perform a simulation study on simulated data, focusing on assessing the Kingman hypothesis, $\Lambda = \delta_0$, which can be robustly evaluated based upon whether or not $\lambda_3 = 0$. The type space consists of 10 binary loci, or 2^{10} types, with mutations flipping a uniformly chosen locus, i.e. the matrix M has 10 identical, non-zero entries in each row corresponding to flipping one locus each. The total mutation rate is set at $\theta = 0.1$. Samples of 20 lineages were simulated at each of five time points from both the Kingman ($\lambda_3 = 0$) and Bolthausen-Sznitman ($\Lambda = U(0, 1)$, $\lambda_3 = 0.5$) coalescents. The data sets are summarised in Table 3. Both data sets come from independent simulations, and are sampled from a population at stationarity. The Kingman coalescent is a classical model of genetic ancestry, while the Bolthausen-Sznitman coalescent has recently been suggested as a ancestral model for influenza and HIV [Neher and Hallatschek, 2013].

We set $\eta = 10^{-6}$ and specify the prior for the density of Λ on $[\eta, 1]$ as a Dirichlet process mixture model of truncated Gaussian kernels. The base measure is the uniform measure on $[\eta, 1]$, with total mass scaled to 0.1. Finally, the prior for $\tau^{-1/2}$, the standard deviations of the truncated Gaussian kernels was chosen to be the Beta(1.0, 3.0) distribution on $[0, 1]$. Truncating the

Time	Bolthausen-Sznitman	Kingman
0.0	20 x 1001001111	20 x 0010000000
0.5	19 x 1001001111	15 x 0010000000
	1 x 1101001111	5 x 0000000000
1.0	20 x 1001001111	8 x 0010000000
		6 x 0000000000
		6 x 0010001000
1.5	19 x 1001001111	10 x 0010000000
	1 x 1001101111	6 x 0000000000
		4 x 0010001000
2.0	19 x 1001001111	16 x 0010000000
	1 x 1001001110	4 x 0010001000

Table 3: Simulated observed sequences from the two models with $\theta = 0.1$ and M as described in the paragraph above.

maximal standard deviation at 1 excludes some very flat densities from the support of the prior, but the standard normal density is already very flat across $[\eta, 1]$ and the truncation was found to yield substantial gains in speed of convergence of algorithms. Note that neither data generating model lies in the support of this prior, but both can be well approximated by members of the support. The choice of hyperparameters was made because it yields a relatively flat marginal prior for λ_3 , the quantity of interest (c.f. Figure 3), and the prior satisfies the requirements of the consistency result in Theorem 2.

We make use of the Sethumaran stick-breaking construction of the Dirichlet process [Sethumaran, 1994] and truncate our prior after the first four atoms. For our choice of base measure and concentration parameter this results in a total variation truncation error of order $400e^{-30} \approx 3.7 \times 10^{-11}$ [Ishwaran and James, 2001, Theorem 2]. Any truncation error could be avoided by pushing forward the prior directly onto the space of moment sequences as illustrated in Section 5. The cost is a more computationally expensive prior to sample and evaluate, as well as a higher dimensional parameter space consisting of 98 moments for our data sets. We do not investigate this strategy further in this paper.

The four atom truncation results in 11 parameters: four locations and standard deviations of truncated Gaussian kernels, and three stick break points. The fourth break point is set to fulfil the constraint of the weights summing to 1. We propose updates to these parameters using a truncated Gaussian random walk on $[\eta, 1]^4 \times [0, 1]^4 \times [0, 1]^3$ with covariance matrix 0.0025Id . This scaling was found to result in a reasonable balance of acceptance probability and jump size for the first moment λ_3 .

We approximate the likelihoods required for computing the acceptance probability a using a straightforward adaptation of the optimised importance sampling method of Koskela et al. [2015a] to the time series setting of [Beaumont, 2003], but need to specify the number of particles to use. More particles will result in more accurate approximations, but at greater computational cost. In [Doucet et al., 2015] the authors show that tuning the variance of the log likelihood estimator to 1.44 results in efficient algorithms under a wide range of assumptions. Preliminary simulations showed this was achieved in our setting by choosing 75 particles for Bolthausen-Sznitman data, and 180 particles for Kingman data.

Remark 7. In the context of real data, when the true data generating parameters are not known, optimising the number of particles using trial runs may require an infeasible amount of computation. In practice, adaptive algorithms, which optimise parameters online, can be used to circumvent this problem. Recent work by Sherlock et al. [2015] has shown how to implement such adaptivity to all four algorithms discussed below, while maintaining ergodicity and the

correct stationary distribution of the algorithm.

It is well known that the standard, exact pseudo-marginal algorithm suffers from “sticking” behaviour, where an unusually high likelihood estimator prevents the algorithm from moving for a macroscopic number of steps [Andrieu and Roberts, 2009]. The usual solution is to use a noisy version of the algorithm, in which the likelihood estimator is recomputed at each stage. This doubles the number of required likelihood evaluations and biases the algorithm into an incorrect stationary distribution, but can greatly reduce the variance of estimates. We compare both the exact and noisy versions of the pseudo-marginal algorithm in Figure 2. We also investigate the effect of delayed acceptance acceleration [Christen and Fox, 2005], in which proposed moves are first subjected to an accept-reject decision based on an approximate likelihood function that is cheap to compute. Only samples which are accepted at this first stage are subjected to an accept-reject decision based on the full likelihood estimates, or more specifically a slight modification to ensure that the delayed acceptance mechanism does not affect the stationary distribution of the algorithm. In the Λ -coalescent setting approximate likelihoods are readily available in the form of Product of Approximate Conditionals (or PAC) methods [Koskela et al., 2015a, Section 4.3], which we use to implement delayed acceptance chains.

Figure 2 shows trace plots of 20 000 steps from the four algorithms introduced above. The exact pseudo-marginal algorithm exhibits sticking behaviour as might be expected, but it is surprising to see that the noisy algorithm does not completely eliminate it. We conjecture that the remaining stickiness in the noisy trace plot is due to multiple, narrow modes in the 11 dimensional posterior. It is also clear that the bias in the noisy algorithm is confounding the signal in the data, as the traces are much more intermixed than those of the exact algorithm.

Both the noisy and exact pseudo-marginal algorithm are very computationally expensive to run, particularly for the Kingman data set due to the larger number of particles used to estimate likelihoods. Delayed acceptance acceleration reduces these run times as expected, particularly for the Kingman case. Both delayed acceptance algorithms also suffer from sticking, and show less clear separation of the traces than the exact algorithm. They also look very similar to each other.

Since it appears to be difficult to eliminate sticking behaviour in this case, we chose to leverage the speed up obtained by making use of delayed acceptance and ran a further exact, delayed acceptance pseudo-marginal algorithm for 200 000 steps. A trace plot is shown in Figure 3. Sticking behaviour is still present, but on a much shorter scale relative to the run length. Run times are comparable to the noisy algorithm without delayed acceptance, and the Bolthausen-Sznitman trace is again clearly centred at a higher level than the Kingman trace. We thinned the output of this long run by a factor of 4 000 to reduce the effect of sticking points to obtain 50 samples of first moments, which were used to plot the histograms in Figure 3.

It is clear from both the trace plots and histograms in Figure 3 that the run length is still not sufficient for fully converged estimates. However, both plots already show a clear shift of posterior modes toward the values generating the data. The red histogram is consistent with the Kingman coalescent, while the blue one is consistent with the Bolthausen-Sznitman coalescent. Moreover, approximate 95% credible intervals are $\lambda_3 \in [0.1, 0.6]$ for the Bolthausen-Sznitman posterior, and $\lambda_3 \in [\eta, 0.5]$ for the Kingman posterior. This suggests the relatively short time series is nevertheless sufficiently informative to reject the incorrect model in both cases.

8 Discussion

In this paper we have presented a robust framework for Bayesian non-parametric inference under Λ -coalescent processes for time series data, and studied the feasibility of implementable

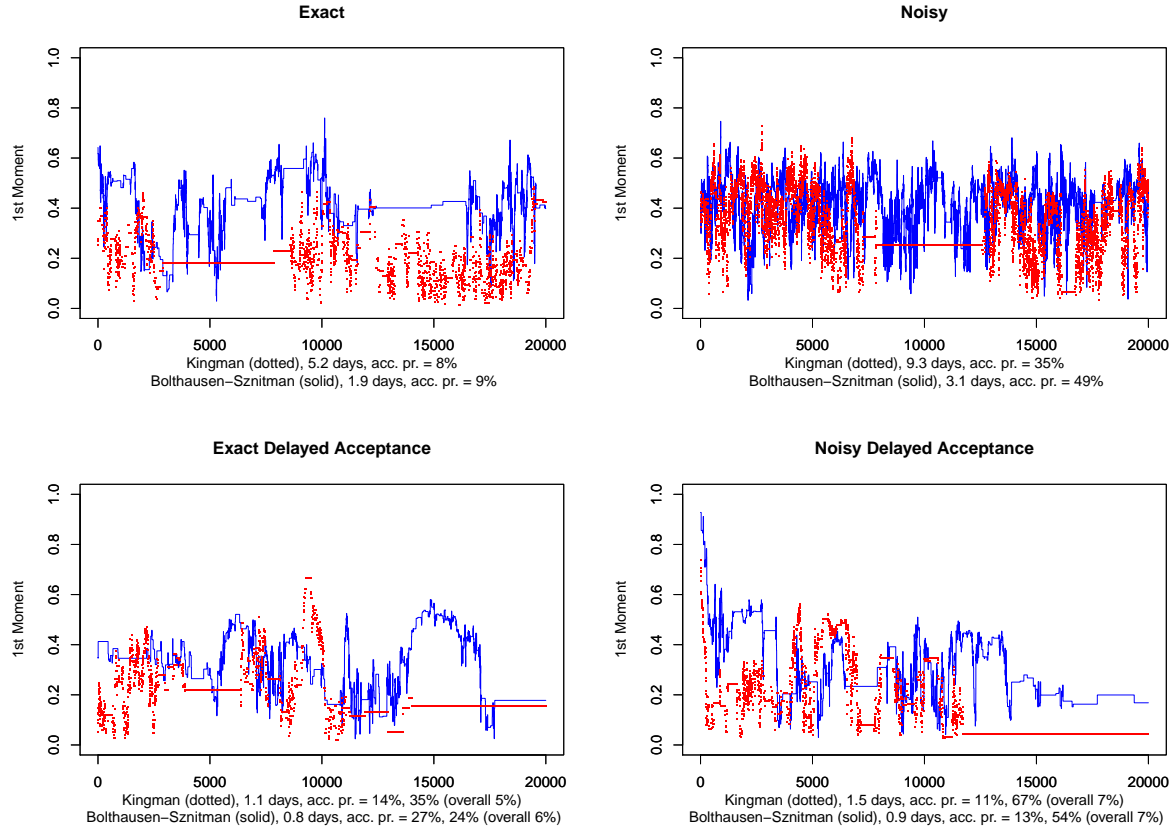


Figure 2: Trace plot of the pseudo-marginal algorithm (top left), the noisy algorithm (top right) and corresponding delayed acceptance algorithms (bottom row). Also shown are computation times (on a mid-range Toshiba laptop with an Intel i5 processor) and acceptance probabilities. Delayed acceptance runs show acceptance probabilities for both stages, as well as an overall probability. All runs are independent and initialised from the prior.




Figure 3: (Left) Trace plot of the delayed acceptance exact pseudo-marginal algorithm, along with computation times (on a mid-range Toshiba laptop with an Intel i5 processor) and acceptance probabilities for both stages as well as overall. Both runs are independent and initialised from the prior. (Right) Histograms of both Kingman (dashed) and Bolthausen-Sznitman (solid) posteriors estimated from 50 MCMC samples obtained by thinning the runs shown on the left. The estimated prior density is shown (dotted), based on 10 000 independent samples from the prior.

families of algorithms for practical inference. We obtained posterior consistency under verifiable conditions on the prior, and demonstrated that time series data is necessary for identifiability of the Λ -measure. As seen in Example 1, lack of consistency can lead to very low statistical power and high sensitivity of inference both to confounding parameters, such as mutation rate, and the observed allele frequencies. A theoretical guarantee of consistency is crucial as expressions for statistical power rely on intractable stationary distributions and transition densities of Λ -Fleming-Viot jump-diffusions, making the reliability of experiments without time series data very difficult to evaluate.

Efficient methods for importance sampling Λ -coalescent trees are available [Birkner et al., 2011, Koskela et al., 2015a], and these can be used to generalise the pseudo-marginal MCMC algorithms of [Beaumont, 2003] for temporally spaced data. The consistency conditions of Theorem 2 on the prior are sufficiently mild to permit the use of Dirichlet process mixture model priors, which can be readily truncated for implementable algorithms. Alternatively, we have shown that parametrising the inference problem via truncated moment sequences leads to implementable algorithms with no discretisation or truncation error. This work provides a strong indication that time series data, and accompanying inference methods such as the one outlined above, should be adopted as standard whenever the coalescent generating the data cannot be assumed to be known.

Generalising of consistency result within the Λ -Fleming-Viot process class to include unknown drift, which can be used to model e.g. mutation, recombination and selection, as well as more general Λ -measures is of great interest. However, it is difficult for a number of reasons. Firstly, relaxing conditions on Λ near 0 while ensuring the integral in (3) remains finite is challenging. Likewise, it is well known that equivalent changes of measure for Lévy processes necessitate equivalent Lévy measures (see e.g. [Sato, 1999], Theorem 33.1), and this is also the condition needed for the jump-diffusions considered in [Cheridito et al., 2005]. The way in which the drift can be transformed while maintaining absolute continuity in [Cheridito et al., 2005] is also restrictive, and depends on the diffusion coefficient and Lévy compensator. Finally, any difference in diffusion coefficients will obviously destroy absolute continuity outright, so if there were an atom $\Lambda(\{0\}) > 0$, its size would have to be known with certainty.

It would also be of great interest to obtain contraction rates of the posterior under verifiable conditions. Obtaining rates is a challenging problem in non-i.i.d. Bayesian non-parametric inference, and existing results by Gugushvili et al. [2015] for compound Poisson processes and Nickl and Söhl [2015] for scalar diffusions do not seem generalisable. A different approach by Nguyen [2013] for mixing measures of infinite mixture models could present a promising directions of future work by viewing the Λ -coalescent tree as a mixture of merger events, but adaptation into the present setting is a formidable task and is beyond the scope of this paper.

The method of parametrising the unknown Λ -measure by its first $n - 2$ moments when the data set is of size $n \in \mathbb{N}$ reflects the limited amount of signal in finite data. More precisely, the likelihood given a sample of size $n \in \mathbb{N}$ is constant within moment classes of order n , so that any variation in the posterior within these moment classes is due solely to the prior. Hence this parametrisation can be seen as regularising an under-determined inference problem in an infinite dimensional space by identifying an appropriate, data-driven, finite dimensional quotient space in which to conduct inference. We believe this approach to have more broad applicability in non-parametric statistics as well as an alternative to direct regularisation by a prior in the infinite dimensional space, or to approximate projections onto finite dimensional subspaces [Cui et al., 2014].

The algorithms used to approximate the posterior and maximise/minimise quantities of interest given the posterior are highly computationally intensive, and we do not expect our approach to be competitive with well-chosen parametric families when the number of observed lineages or loci

is large. However, the simulations in Section 7 demonstrate that our assumption-free framework can be used to empirically evaluate the modelling fit of parametric families given moderately sized pilot data, for instance by ensuring that the family contains a candidate Λ which matches the MAP estimators of some small number of moments. Such parametric families can then be confidently used to process larger data sets. The pseudo-marginal method can also be adapted to incorporate unknown mutation parameters, recombination and other forces not considered in this paper, albeit at the cost of greater computational cost and lower parameter identifiability. This cost can be alleviated to a large extent by modern GPU and cluster computing approaches, because the importance sampling algorithm used to estimate likelihoods is readily parallelisable. For example, up to 500 fold speed up was reported by Lee et al. [2010] when computations were parallelised on GPUs instead of being run in serial on CPUs. Such gains in computation speed would make the algorithms employed in Section 7 practical for many realistic genetic data sets.

Acknowledgements

The authors are grateful to Tim Sullivan for insight into orthogonal polynomials as well as Bayesian brittleness, to Felipe Medina Aguayo for fruitful conversations about pseudo-marginal methods, and to Yun Song for helpful comments on Remark 2. Jere Koskela is supported by EPSRC as part of the MASDOC DTC at the University of Warwick. Grant No. EP/HO23364/1. Paul Jenkins is supported in part by EPSRC grant EP/L018497/1.

References

- E. C. Anderson. An efficient Monte Carlo method for estimating N_e from temporally spaced samples using a coalescent-based likelihood. *Genetics*, 170(2):955–967, 2005.
- C. Andrieu and G. O. Roberts. The pseudo-marginal approach for efficient Monte Carlo computations. *Ann. Stat.*, 37(2):697–725, 2009.
- E. Árnason. Mitochondrial cytochrome b DNA variation in the high-fecundity Atlantic cod: trans-Atlantic clines and shallow gene genealogy. *Genetics*, 166:1871–1885, 2004.
- M. A. Beaumont. Estimation of population growth or decline in genetically monitored populations. *Genetics*, 164:1139–1160, 2003.
- J. Berestycki, N. Berestycki, and J. Schweinsberg. The genealogy of branching Brownian motion with absorption. *Ann. Probab.*, 41(2):527–618, 2013.
- J. Bertoin and J.-F. Le Gall. Stochastic flows associated to coalescent processes. *Probab. Theory Related Fields*, 126:261–288, 2003.
- J. Bertoin and J.-F. Le Gall. Stochastic flows associated to coalescent processes II: stochastic differential equations. *Ann. Inst. Henri Poincaré Probab. Stat.*, 41(3):307–333, 2005.
- A. Bhattacharya and D. B. Dunson. Strong consistency of nonparametric Bayes density estimation on compact metric spaces with applications to specific manifolds. *Ann. Inst. Stat. Math.*, 64:687–714, 2012.
- M. Birkner and J. Blath. Computing likelihoods for coalescents with multiple collisions in the infinitely many sites model. *J. Math. Biol.*, 57(3):435–463, 2008.
- M. Birkner and J. Blath. Measure-valued diffusions, general coalescents and population genetic inference. in J. Blath, P. Mörters, M. Scheutzow (Eds.), *Trends in Stochastic Analysis*, LMS 351:329–363, 2009.

- M. Birkner, J. Blath, and M. Steinrück. Importance sampling for Lambda-coalescents in the infinitely many sites model. *Theor. Popln Biol.*, 79(4):155–173, 2011.
- J. P. Bollback, T. L. York, and R. Nielsen. Estimation of $2N_e s$ from temporal allele frequency data. *Genetics*, 179(1):497–502, 2008.
- J. D. G. Boom, E. G. Boulding, and A. T. Beckenback. Mitochondrial DNA variation in introduced populations of Pacific oyster, *Crassostrea gigas*, in British Columbia. *Can. J. Fish. Aquat. Sci.*, 51:1608–1614, 1994.
- P. Cheridito, D. Filipović, and M. Yor. Equivalent and absolutely continuous measure changes for jump-diffusion processes. *Ann. Appl. Probab.*, 15(3):1713–1732, 2005.
- J. A. Christen and C. Fox. Markov chain Monte Carlo using an approximation. *J. Comput. Graph. Stat.*, 14(4):795–810, 2005.
- T. Cui, J. Martin, Y. M. Marzouk, A. Solonen, and A. Spantini. Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems*, 30(11):114015, 2014.
- R. Der and J. B. Plotkin. The equilibrium allele frequency distribution for a population with reproductive skew. *Genetics*, 196(4):1199–1216, 2014.
- P. Diaconis and D. Freedman. On the consistency of Bayes estimates. *Ann. Statist.*, 14(1):1–26, 1986.
- P. Donnelly and T. Kurtz. Particle representations for measure-valued population models. *Ann. Probab.*, 27(1):166–205, 1999.
- A. Doucet, M. Pitt, G. Deligiannidis, and R. Kohn. Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika*, 102(2):295–313, 2015.
- A. J. Drummond, G. K. Nicholls, A. G. Rodrigo, and W. Solomon. Estimating mutation parameters, population history and genealogy simultaneously from temporally spaced sequence data. *Genetics*, 161(3):1307–1320, 2002.
- A. J. Drummond, A. Rambaut, B. Shapiro, and O. G. Pybus. Bayesian coalescent inference of past population dynamics from molecular sequences. *Mol. Biol. Evol.*, 22:1185–1192, 2005.
- R. Durrett and J. Schweinsberg. A coalescent model for the effect of advantageous mutations on the genealogy of a population. *Stoch. Proc. Appl.*, 115:1628–1657, 2005.
- B. Eldon and J. Wakeley. Coalescent processes when the distribution of offspring number among individuals is highly skewed. *Genetics*, 172:2621–2633, 2006.
- T. S. Ferguson. A Bayesian analysis of some nonparametric problems. *Ann. Stat.*, 1(2):209–230, 1973.
- S. Gugushvili, F. van der Meulen, and P. Spreij. Nonparametric Bayesian inference for multidimensional compound Poisson processes. *Mod. Stoch. Theory Appl.*, 2(1):1–15, 2015.
- J. Gurland. Inversion formulae for the distribution of ratios. *Ann. Math. Statist.*, 19:228–237, 1948.
- R. C. Hannum, M. Hollander, and N. A. Langberg. Distributional results for random functionals of a Dirichlet process. *Ann. Probab.*, 9:665–670, 1981.
- N. L. Hjort, C. Holmes, P. Müller, and S. G. Walker, editors. *Bayesian nonparametrics*. Cambridge series in statistical and probabilistic mathematics. Cambridge University Press, 2010.
- M. Hoffmann and A. Olivier. Nonparametric estimation of the division rate of an age dependent branching process. *Stoch. Proc. Appl.*, 126(5):1433–1471, 2016.

- R. Höpfner, M. Hoffmann, and E. Löcherbach. Nonparametric estimation of the death rate in branching diffusions. *Scand. J. Stat.*, 29:665–690, 2002.
- H. Ishwaran and L. F. James. Gibbs sampling methods for stick-breaking priors. *J. Amer. Statist. Assoc.*, 96(453):161–173, 2001.
- J. F. C. Kingman. The coalescent. *Stochast. Process. Applic.*, 13(3):235–248, 1982.
- J. Koskela, P. A. Jenkins, and D. Spanò. Computational inference beyond Kingman’s coalescent. *J. Appl. Probab.*, 52(2):519–537, 2015a.
- J. Koskela, D. Spanò, and P. A. Jenkins. Consistency of Bayesian nonparametric inference for discretely observed jump diffusions. *Preprint*, arXiv:1506.04709, 2015b.
- A. Lee, C. Yau, M. B. Giles, A. Doucet, and C. C. Holmes. On the utility of graphics cards to perform massively parallel simulation of advanced Monte Carlo methods. *J. Comp. Graph. Stat.*, 19(4):769–789, 2010.
- A. Y. Lo. On a class of Bayesian nonparametric estimates. 1. density estimates. *Ann. Statist.*, 12:351–357, 1984.
- A.-S. Malaspinas, O. Malaspinas, S. N. Evans, and M. Slatkin. Estimating allele age and selection coefficient from time-serial data. *Genetics*, 192(2):599–607, 2012.
- I. Mathieson and G. McVean. Estimating selection coefficients in spatially structured populations from time series data of allele frequencies. *Genetics*, 193(3):973–984, 2013.
- V. N. Minin, E. W. Bloomquist, and M. A. Suchard. Smooth skyride through a rough skyline: Bayesian coalescent-based inference of population dynamics. *Mol. Biol. Evol.*, 25:1459–1471, 2008.
- M. Möhle. The concept of duality and applications to Markov processes arising in neutral population genetics models. *Bernoulli*, 5(5):761–777, 1999.
- M. Möhle. On sampling distributions for coalescent processes with simultaneous multiple collisions. *Bernoulli*, 12(1):35–53, 2006.
- R. A. Neher and O. Hallatschek. Genealogies of rapidly adapting populations. *Proc. Natl Acad. Sci.*, 110(2):437–442, 2013.
- X. Nguyen. Convergence of latent mixing measures in finite and infinite mixture models. *Ann. Stat.*, 41(1):370–400, 2013.
- R. Nickl and J. Söhl. Nonparametric Bayesian posterior contraction rates for discretely observed scalar diffusions. *Preprint*, arXiv:1510.05526, 2015.
- H. Owhadi, C. Scovel, and T. Sullivan. Brittleness of Bayesian inference under finite information in a continuous world. *Electron. J. Statist.*, 9(1):1–79, 2015.
- J. Pitman. Coalescents with multiple collisions. *Ann. Probab.*, 27(4):1870–1902, 1999.
- E. Regazzini, A. Guglielmi, and G. Di Nunno. Theory and numerical analysis for exact distributions of functionals of a Dirichlet process. *Ann. Stat.*, 30(5):1376–1411, 2002.
- S. Sagitov. The general coalescent with asynchronous mergers of ancestral lineages. *J. Appl. Probab.*, 36(4):1116–1125, 1999.
- K.-I. Sato. *Lévy processes and infinitely divisible distributions*. Cambridge University Press, 1999.
- J. Schweinsberg. Rigorous results for a population model with selection II: genealogy of the population. *Preprint*, arXiv:1507.00394, 2015.

- J. Sethuraman. A constructive definition of Dirichlet priors. *Stat. Sinica*, 4:639–650, 1994.
- C. Sherlock, A. Golightly, and D. A. Henderson. Adaptive, delayed-acceptance MCMC for targets with expensive likelihoods. *Preprint*, arXiv:1509.00172, 2015.
- J. P. Spence, J. A. Kamm, and Y. S. Song. The site frequency spectrum for general coalescents. *Genetics*, 202:1549–1561, 2016.
- M. Steinrücken, M. Birkner, and J. Blath. Analysis of DNA sequence variation within marine species using Beta-coalescents. *Theor. Popln Biol.*, 87:15–24, 2013.
- F. van der Meulen and H. van Zanten. Consistent nonparametric Bayesian inference for discretely observed scalar diffusions. *Bernoulli*, 19(1):44–63, 2013.
- J. Wang. Regularity of semigroups generated by Lévy type operators via coupling. *Stoch. Proc. Appl.*, 120(9):1680–1700, 2010.
- G. Winkler. Extreme points of moment sets. *Math. Oper. Res.*, 30(4):581–587, 1988.